arXiv:1703.08570v1 [math.OC] 24 Mar 2017

Stochastic Methods for Composite Optimization Problems

John C. Duchi[1,2] and Feng Ruan[2]

{jduchi,fengruan}@stanford.edu

Departments of [1]Electrical Engineering and [2]Statistics
Stanford University

Abstract

We consider minimization of stochastic functionals that are compositions of a (potentially) non-smooth convex function h and smooth function c. We develop two stochastic methods—a stochastic prox-linear algorithm and a stochastic (generalized) sub-gradient procedure—and prove that, under mild technical conditions, each converges to first-order stationary points of the stochastic objective. We provide experiments further investigating our methods on non-smooth phase retrieval problems; the experiments indicate the practical effectiveness of the procedures.

1 Introduction

Let $f : \mathbb{R}^d \to \mathbb{R}$ be the stochastic composite function

$$f(x) := \mathbb{E}_P[h(c(x; S); S)] = \int_{\mathcal{S}} h(c(x; s); s) dP(s),$$

where P is a probability distribution on a sample space $\mathcal{S}$ and for each $s \in \mathcal{S}$, the function $z \mapsto h(z; s)$ is closed convex and $x \mapsto c(x; s)$ is smooth. In this paper, we consider stochastic methods for minimization—or at least finding stationary points—of such composite functionals, studying the problem

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & f(x) + \varphi(x) \\ \text{subject to} \quad & x \in X, \end{aligned} \tag{1}$$

where $X \subset \mathbb{R}^d$ is a closed convex set and $\varphi : \mathbb{R}^d \to \mathbb{R}$ is a closed convex function.

A number of problems are representable in the form (1). Of course, taking the function c as the identity mapping, classical regularized stochastic convex optimization problems fall into this framework [24], including regularized least-squares and the Lasso [18, 32], with $s = (a, b) \in \mathbb{R}^d \times \mathbb{R}$ and $h(x; s) = \frac{1}{2}(a^T x - b)^2$ and φ typically some norm on x; the support vector machine problem [9], with $s = (a, b) \in \mathbb{R}^d \times \{-1, 1\}$ and $h(x; s) = \left[1 - ba^T x\right]_+$. The more general problem (1) includes a number of important non-convex problems. Examples include non-linear least squares [cf. 26], with $s = (a, b)$ and $b \in \mathbb{R}$, the convex term $h(t; s) \equiv h(t) = \frac{1}{2}t^2$ is independent of the sampled s, and $c(x; s) = c_0(x; a) - b$ where c_0 is some smooth function a modeler believes predicts b well given $x \in \mathbb{R}^d$ and data a. Another compelling example is the (robust) phase retrieval problem [6, 30]—which we explore somewhat more in depth in our numerical experiments—where the data $s = (a, b) \in \mathbb{R}^d \times \mathbb{R}_+$, $h(t; s) \equiv h(t) = |t|$ or $h(t; s) \equiv h(t) = \frac{1}{2}t^2$, and $c(x; s) = (a^T x)^2 - b$. In the case that $h(t) = |t|$, the form (1) is a natural exact penalty for the solution of a collection of quadratic equalities $(a_i^T x)^2 - b_i = 0$, $i = 1, \ldots, N$, where we take P to be point masses on pairs (a_i, b_i).

Fletcher and Watson [17, 16] initiated study of the non-stochastic version of the composite problem (1), that is,

$$\underset{x}{\text{minimize}} \quad h(c(x)) + \varphi(x), \;\; \text{subject to} \;\; x \in X \tag{2}$$

for fixed convex h, smooth c, convex φ and convex X. A substantial motivation of this early work is nonlinear programming problems with the equality constraint that $x \in \{x : c(x) = 0\}$, in which case taking $h(z) = \|z\|$ for some norm $\|\cdot\|$ functions as an exact penalty [19] for the equality constraint $c(x) = 0$. A more recent line of work, beginning with Burke [5] and continued (variously) by Druvyatskiy, Ioffe, Kempton, Lewis, and Wright [22, 13, 12, 11], establishes first-order convergence guarantees, as well as rates of convergence, for iterative methods that solve sequentially constructed (local) convex surrogates for the problem (2).

Roughly, these papers construct a model of the composite function $f(x) = h(c(x))$ as follows. Letting $\nabla c(x)$ be the transpose of the Jacobian of c at x, so $c(y) = c(x)+\nabla c(x)^T(y-x)+o(\|y-x\|)$, one defines the "linearized" model of f at x by

$$f_x(y) := h(c(x) + \nabla c(x)^T(y - x)), \tag{3}$$

which is evidently convex in y. When h and ∇c are Lipschitzian, then evidently $|f_x(y) - f(x)| = O(\|x - y\|^2)$, so that the model (3) is second-order accurate, which motivates the following *prox-linear* method. Beginning from some $x_0 \in X$, iteratively construct x_k via

$$x_{k+1} = \operatorname*{argmin}_{x \in X} \left\{ f_{x_k}(x) + \varphi(x) + \frac{1}{2\alpha_k} \|x - x_k\|^2 \right\}, \tag{4}$$

where $\alpha_k > 0$ is a stepsize that may be chosen by a line-search. For small α_k, the iterates (4) guarantee decreasing $h(c(x_k)) + \varphi(x_k)$, the sequence of problems (4) are convex, and moreover, the iterates x_k converge to stationary points of problem (2) [12, § 5]. The prox-linear method is effective so long as minimizing the models $f_{x_k}(x)$ is reasonably computationally easy.

In our stochastic composite problem (1) where $f(x) = \mathbb{E}[h(c(x; S); S)]$, the iterates (4) may be computationally challenging. Even in the case in which P is discrete so that problem (1) has the form $f(x) = \frac{1}{n}\sum_{i=1}^n h_i(c_i(x))$, which is evidently of the form (2), the iterations generating x_k may be prohibitively expensive for large n. When P is not discrete or when it is unknown, because we can only simulate draws $S \sim P$ or in natural statistical settings in which the only access to P is via observations $S_i \stackrel{\text{iid}}{\sim} P$, then the iteration (4) is essentially infeasible. Given the wide applicability of the stochastic composite problem (1), however, it is of substantial interest to develop efficient online and stochastic methods to (approximately) solve it, or at least to find local optima.

In this work, we develop and study two stochastic algorithms: a stochastic linear proximal algorithm, which is a stochastic analogue of problem (4), and a stochastic subgradient algorithm, both of whose definitions we give in Section 2. An advantage of these methods is that their iterations are often computationally simple, and they require only individual samples $S \sim P$ at each iteration. Consider for concreteness the case when P is discrete and supported on $i = 1, \ldots, n$ (i.e. $f(x) = \frac{1}{n}\sum_{i=1}^n h_i(c_i(x))$). Then instead of solving the non-trivial subproblem (4), the stochastic prox-linear algorithm samples $i_0 \in [n]$ uniformly, then substitutes h_{i_0} and c_{i_0} for h and c in the iteration. Thus, as long as there is a prox-linear step for the individual compositions $h_i \circ c_i$, the algorithm is reasonably easy to implement and execute.

The main result of this paper is that the stochastic prox-linear and subgradient methods we develop for the composite optimization problem are convergent. More precisely, under a few mild technical conditions, both the stochastic prox-linear and subgradient iterations converge to stationary points of the (potentially) non-smooth, non-convex objective (1) (Theorem 1 in Sec. 2 and Theorem 5 in Sec. 3.4). As gradients $\nabla f(x)$ may not exist (and may not even be zero at stationary points because of the non-smoothness of the objective), demonstrating this convergence provides some challenge. To circumvent these difficulties, we show that the iterates are asymptotically equivalent to the trajectories of a particular ordinary differential inclusion [1] (a non-smooth

generalization of ordinary differential equations (ODEs)) related to problem (1), building off of the classical ODE method [23, 21, 4] (see Section 3.2). By developing a number of analytic properties of the limiting differential inclusion using the composite structure $h \circ c$, we show that trajectories of the ODE must converge (Section 3.3). A careful stability analysis then shows that limit properties of trajectories of the ODE are preserved under small perturbations, and viewing our algorithms as noisy discrete approximations to a solution of the ordinary differential inclusion gives our desired convergence (Section 3.4).

Our results do not, as yet, provide rates of convergence for the stochastic procedures, so to investigate the properties of the methods we propose, we perform a number of numerical simulations in Section 4. We focus on a discrete version of problem (1) with the robust phase retrieval objective $f(x; a, b) = |(a^T x)^2 - b|$, which facilitates comparison with deterministic methods (4). Our experiments corroborate our theoretical predictions, showing the advantages of stochastic over deterministic procedures for even some medium-scale problems, and they also show that the stochastic prox-linear method may be preferable to stochastic subgradient methods because of nice robustness properties it enjoys (which our simulations verify, though our theory does not yet explain).

Notation and basic definitions We collect here our (mostly standard) notation and basic definitions that we require. We let $\mathbb{B}$ denote the unit ℓ_2-ball in $\mathbb{R}^d$, where the dimension d is apparent from context, and $\|\cdot\|$ denotes the standard Euclidean norm. For a function $f : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$, we let $\partial f(x)$ denote the Fréchet subdifferential (also called the regular subdifferential [29, Ch. 8.B]) of f at the point x, which is defined as

$$\partial f(x) := \left\{ g \in \mathbb{R}^d : f(y) \geq f(x) + \langle g, y - x \rangle + o(\|y - x\|) \text{ as } y \to x \right\}.$$

The Fréchet subdifferential and subdifferential coincide for convex functions [29, Ch. 8]. We define the (Clarke) directional derivative of a function f at the point x in direction v by

$$f'(x; v) := \liminf_{t \downarrow 0, v' \to v} \frac{f(x + tv) - f(x)}{t},$$

and recall [29, Ex. 8.4] that $\partial f(x) = \{w \in \mathbb{R}^d : \langle v, w \rangle \leq f'(x; v) \text{ for all } v\}$.

We let $\mathcal{C}(A, B)$ denote the continuous functions from the set A to the set B. Given a sequence of functions $f_n : \mathbb{R}_+ \to \mathbb{R}^d$, we say that $f_n \to f$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ if $f_n \to f$ uniformly on all compact sets, that is, for all $T < \infty$ we have

$$\lim_{n \to \infty} \sup_{t \in [0,T]} \|f_n(t) - f(t)\| = 0.$$

This is equivalent to convergence in the metric $\mathsf{d}(f, g) := \sum_{t=1}^{\infty} 2^{-t} \sup_{\tau \in [0,t]} \|f(\tau) - g(\tau)\| \wedge 1$, which shows the standard result that $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ is a Fréchet space. For a closed convex set X, we let $\mathbb{I}_X$ denote the $+\infty$-valued indicator for X, that is, $\mathbb{I}_X(x) = 0$ if $x \in X$ and $+\infty$ otherwise. The normal cone to X at x is

$$\mathcal{N}_X(x) := \{v \in \mathbb{R}^d : \langle v, y - x \rangle \leq 0 \text{ for all } y \in X\}.$$

For closed convex sets C, $\pi_C(x) := \operatorname{argmin}_{y \in C} \|y - x\|$ denotes the Euclidean projection of x onto C. For a matrix A, we let $|\!|\!| A |\!|\!|_{\text{op}} := \sup_{\|u\|=1} \|Au\|$ be the ℓ_2-operator norm.

2 Algorithms and Main Convergence Result

In this section, we introduce two natural algorithms for problem (1), which we call the stochastic prox-linear and subgradient methods. The first is a generalization of the prox-linear method Burke [5] develops, whose analytic and other properties have been further investigated by Drusvyatskiy, Ioffe, Kempton, and Lewis [12, 13, 11]. The second is the natural generalization of the simple subgradient descent method [15].

We begin with the stochastic linear proximal method. For this method, we require a particular "linearization" of the instantaneous objective $h(c(x;s);s)$, where we linearize the internal function c without linearizing h. To that end, we define the function

$$f_x(y;s) := h(c(x;s) + \nabla c(x;s)^T(y-x);s)$$

where $\nabla c(x;s) \in \mathbb{R}^{m\times d}$ is the gradient matrix of the function $c(\cdot;s)$ at the point x. The stochastic prox-linear method is then

$$\begin{aligned} &\text{Draw} \quad S_k \overset{\text{iid}}{\sim} P \\ &x_{k+1} := \operatorname*{argmin}_{y\in X} \left\{ f_{x_k}(y;s_k) + \varphi(y) + \frac{1}{2\alpha_k} \left\|y - x_k\right\|^2 \right\}. \end{aligned} \tag{5}$$

We consider a variant method that is in many cases even simpler to implement. In particular, let $\mathsf{g}(x;s) \in \nabla c(x;s)\partial h(c(x;s);s)$ be a (fixed, as chosen by the subgradient oracle conditional on s) element of the Fréchet subdifferential of $h(c(x;s);s)$. Then the stochastic projected subgradient algorithm for problem (1) is

$$\begin{aligned} &\text{Draw} \quad S_k \overset{\text{iid}}{\sim} P \quad \text{and set } g_k = \mathsf{g}(x_k;S_k) \\ &x_{k+1} := \operatorname*{argmin}_{y\in X} \left\{ \langle g_k, y\rangle + \varphi(y) + \frac{1}{2\alpha_k} \left\|y - x_k\right\|^2 \right\}. \end{aligned} \tag{6}$$

We choose our sequence of strictly positive stepsizes $\{\alpha_n\}_{k=1}^\infty$ to be square summable but not summable:

$$\sum_{k=1}^\infty \alpha_k = \infty \quad \text{and} \quad \sum_{k=1}^\infty \alpha_k^2 < \infty. \tag{7}$$

For instance, one can choose $\alpha_k \propto k^{-\beta}$, for any $\beta \in (1/2, 1]$.

The main theoretical results of this paper is to show that the above two stochastic algorithms converge almost surely to the stationary points of the objective function $F(x) = f(x) + \varphi(x)$. To state our results formally, we require a few assumptions on the smoothness and continuity properties of the composition $f(x;s) = h(c(x;s);s)$ and the domain X. In particular, we assume that $h(\cdot;s)$ is Lipschitz continuous on appropriate subsets of its domain and that c is smooth, that is, $\nabla c(\cdot;s)$ is Lipschitz. Concretely, we define functions $\gamma : \mathbb{R}_+ \times \mathcal{S} \to \mathbb{R}_+$ and $\beta : \mathbb{R}_+ \times \mathcal{S} \to \mathbb{R}_+$ governing the Lipschitzian properties of h and c as follows. For any $B \geq 0$, we assume that on the set $X \cap \{x : \|x\| \leq B\}$ the function $c(\cdot;s)$ has $\beta(B,s)$-Lipschitz gradient, that is,

$$\left\|\nabla c(x;s) - \nabla c(y;s)\right\|_{\text{op}} \leq \beta(B,s) \left\|x-y\right\| \quad \text{for } x, y \in X \cap B\mathbb{B}.$$

We also require $h(\cdot;s)$ to be (locally) Lipschitz in an appropriate sense, specifically, on the domain of possible linearizations $y \mapsto c(x;s) + \nabla c(x;s)^T(y-x)$. That is, we assume that $h(\cdot;s)$ is $\gamma(B,s)$-Lipschitz on the convex set

$$\operatorname{Conv} \bigcup_{x,w} \left\{ c(x;s) + \nabla c(x;s)^T w : x \in X, \|x\| \leq B, \|w\| \leq 1 \right\}.$$

To guarantee well-behavedness of the algorithm, we require the following

Assumption A. *For all $B < \infty$, the Lipschitz constants $\beta(B, s)$ and $\gamma(B, s)$ satisfy*

$$\mathbb{E}[\beta(B, S)^2] < \infty \quad \text{and} \quad \mathbb{E}[\gamma(B, S)^2] < \infty,$$

and there exists $x_0 \in X$ with $\mathbb{E}[\|\nabla c(x; S)\|_{\text{op}}^2] < \infty$.

A simpler version of this assumption is just that $h(\cdot; s)$ is $\gamma(s)$-Lipschitz continuous, but we wish to allow functions h that may grow more quickly than linearly, such as quadratics. Assumption A, as we see later, is sufficient to guarantee that the Fréchet subgradient $\partial f(x)$ exists and is non-empty for all $x \in X$ and is also outer semi-continuous.

With Assumption A in place, we can now proceed to a (mildly) simplified version of our main result in this paper. We let $X^\star$ denote the set of stationary points for the objective function $F(x) = f(x) + \varphi(x)$ over X, that is,

$$X^\star := \{x \in X : \exists g \in \partial f(x) + \partial\varphi(x) \text{ with } \langle g, y - x\rangle \geq 0 \text{ for all } y \in X\}. \tag{8}$$

Equivalently, $\partial f(x) + \partial\varphi(x) \cap -\mathcal{N}_X(x) \neq \emptyset$, or $0 \in \partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x)$. We require one additional assumption for essentially purely technical reasons, which is that the *image* of the stationary points be countable.

Assumption B. *The image $F(X^\star) := \{f(x) + \varphi(x) : x \in X^\star\}$ is countable, that is, $F = f + \varphi$ takes on only countably many values over $X^\star$.*

Of course, if f is convex then $(f + \varphi)(X^\star)$ is a singleton. Moreover, if the set of stationary points $X^\star$ consists of a (finite or countable) collection of sets $X_1^\star, X_2^\star, \ldots$ such that $f + \varphi$ is constant on each $X_i^\star$, then Assumption B holds.

We then have the following convergence result, which is a simplification of our main convergence result, Theorem 5, which we present in Section 3.4.

Theorem 1. *Let Assumptions A and B hold, and assume that X is compact. Let x_k be generated by either of the updates (5) or (6). Then with probability 1, all cluster points of the sequence $\{x_k\}_{k=1}^\infty$ belong to the stationary set $X^\star$ and $F(x_k) = f(x_k) + \varphi(x_k)$ converges.*

3 Convergence Analysis of the Algorithm

In this section, we present the arguments necessary to prove Theorem 1 and its extensions, beginning with a heuristic explanation that we make rigorous subsequently. By inspection and a strong faith in the limiting behavior of random iterations, we might expect that asymptotically the update schemes (5) and (6), as the stepsize $\alpha_k \to 0$, are asymptotically approximately equivalent to iterations of the form

$$\frac{1}{\alpha_k}(x_{k+1} - x_k) \approx -\left[\mathsf{g}(x_k) + v_k + w_k\right] \quad \text{where } \mathsf{g}(x_k) \in \partial f(x_k),\ v_k \in \partial\varphi(x_{k+1}),\ w_k \in \mathcal{N}_X(x_{k+1}),$$

and the correction w_k serves to enforce $x_{k+1} \in X$. As $k \to \infty$ and $\alpha_k \to 0$, we may (again, deferring rigor) treat $\lim_k \frac{1}{\alpha_k}(x_{k+1} - x_k)$ as a continuous time process, and we expect further that the update schemes (5) and (6) are asymptotically equivalent to a continuous time process $t \mapsto x(t) \in \mathbb{R}^d$ that satisfies the differential inclusion (a set-valued generalization of an ordinary differential equation)

$$\begin{aligned} \dot{x} &\in -\partial f(x) - \partial\varphi(x) - \mathcal{N}_X(x) \\ &= -\int \nabla c(x; s)\partial h(c(x; s); s) dP(s) - \partial\varphi(x) - \mathcal{N}_X(x). \end{aligned} \tag{9}$$

We develop a general convergence result showing that this limiting equivalence is indeed the case and that the equality moving from the first to the second line of expression (9) holds. As part of this, we explore in the coming sections how the composite structure $h \circ c$—the convexity of h and smoothness of c—guarantees that the differential inclusion (9) is well-behaved. We begin in Section 3.1 with preliminaries on set-valued analysis and differential inclusions that are necessary for our convergence guarantees, which build on standard convergence results for differential inclusions [1, 20]. Once we have presented these main preliminary results, we show how the stochastic iterations (5) and (6) eventually approximate solution paths to differential inclusions (Section 3.2), which builds off of a number of stochastic approximation results and the so-called "ODE method" as developed by Ljung [23], further studied and extended to differential inclusions by a number of authors (see, for example, the references [21, 2, 4]). We develop the analytic properties of the composite objective, which yields the uniqueness of trajectories solving (9) as well as a particular Lyapunov convergence inequality (Section 3.3). Finally, we develop stability results on the differential inclusion (9), which allows us to prove convergence as in Theorem 1 (Section 3.4).

3.1 Preliminaries: differential inclusions and set-valued analysis

We now review a few results in set-valued analysis and differential inclusions [1, 20]. Our notation and definitions follow closely the standard references of Rockafellar and Wets [29] and Aubin and Cellina [1], and we cite a few results from the book of Kunze [20].

Given a sequence of sets $A_n \subset \mathbb{R}^d$, we define the limit supremum of the sets by limit points of subsequences $y_{n_k} \in A_{n_k}$, that is,

$$\limsup_n A_n := \{y : \exists\ y_{n_k} \in A_{n_k} \text{ s.t. } y_{n_k} \to y \text{ as } k \to \infty\}.$$

We let $G : X \rightrightarrows \mathbb{R}^d$ denote a set-valued mapping G from X to $\mathbb{R}^d$, and we define $\operatorname{dom} G := \{x : G(x) \neq \emptyset\}$. Then G is *outer semicontinuous (o.s.c.)* if for any sequence $x_n \to x \in \operatorname{dom} G$, we have $\limsup_n G(x_n) \subset G(x)$. One says that G is *ϵ-δ outer semicontinuous* [1, Def. 1.1.5] if for all x and $\epsilon > 0$, there exists $\delta > 0$ such that $G(x + \delta\mathbb{B}) \subset G(x) + \epsilon\mathbb{B}$. These notions coincide when $G(x)$ is bounded. Two standard examples of outer-semicontinuous mappings follow.

Lemma 3.1 (Hiriart-Urruty and Lemaréchal [19], Theorem VI.6.2.4)**.** *Let $f : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ be convex. Then the subgradient mapping $\partial f : \operatorname{int} \operatorname{dom} f \rightrightarrows \mathbb{R}^d$ is o.s.c.*

Lemma 3.2 (Rockafellar and Wets [29], Proposition 6.6)**.** *Let X be a closed convex set. Then the normal cone mapping $\mathcal{N}_X : X \rightrightarrows \mathbb{R}^d$ is o.s.c. on X.*

The *differential inclusion* associated with G beginning from the point x_0, denoted

$$\dot{x} \in G(x), \quad x(0) = x_0 \tag{10}$$

has a solution if there exists an absolutely continuous function $x : \mathbb{R}_+ \to \mathbb{R}^d$ satisfying $\frac{d}{dt}x(t) = \dot{x}(t) \in G(x(t))$ for all $t \geq 0$. For $G : \mathcal{T} \rightrightarrows \mathbb{R}^d$ and a measure μ on $\mathcal{T}$, the integral $\int G d\mu$ is

$$\int_{\mathcal{T}} G d\mu = \int_{\mathcal{T}} G(t) d\mu(t) := \left\{ \int g(t) d\mu(t) \mid g(t) \in G(t), g \text{ measurable} \right\}.$$

An outer semicontinuous mapping G is *locally compact* if for all x, the projection of 0 onto $G(y)$, $\pi_{G(y)}(0)$, takes values in some compact set for all y in a neighborhood of x. With these definitions, the following results (with minor extension) on the existence and uniqueness of solutions to differential inclusions are standard.

Lemma 3.3 (Aubin and Cellina [1], Theorem 2.1.4)**.** *Let $G : X \rightrightarrows \mathbb{R}^d$ be outer semicontinuous and compact-valued, and $x_0 \in X$. Assume that there is a compact set $K \subset \mathbb{R}^d$ such that $\pi_{G(x)}(0) \in K$ for all x. Then there exists an absolutely continuous function $x : \mathbb{R}_+ \to \mathbb{R}^d$ such that $\dot{x}(t) \in G(x(t))$ and $x(t) \in x_0 + \int_0^t G(x(\tau))d\tau$ for all $t \in \mathbb{R}_+$.*

Lemma 3.4 (Kunze [20], Theorem 2.2.2)**.** *In addition to the conditions of Lemma 3.3, assume that there exists $c < \infty$ such that*

$$\langle x_1 - x_2, g_1 - g_2 \rangle \leq c \left\| x_1 - x_2 \right\|^2 \quad \text{for } g_i \in G(x_i) \text{ and all } x_i \in \operatorname{dom} G.$$

Then the solution to the differential inclusion (10) *is unique.*

As our final preliminary result, we recall basic Lyapunov theory for differential inclusions. Let $V : X \to \mathbb{R}_+$ be a non-negative function and $W : X \times \mathbb{R}^d \to \mathbb{R}_+$ be continuous and satisfy that $v \mapsto W(x, v)$ is convex for all x. A trajectory $\dot{x} \in G(x)$ is *monotone* for the pair V, W if

$$V(x(T)) - V(x(0)) + \int_0^T W(x(t), \dot{x}(t))dt \leq 0 \quad \text{for } T \geq 0.$$

The following lemma presents sufficient conditions for the existence of such monotone trajectories.

Lemma 3.5 (Aubin and Cellina [1], Theorem 6.3.1)**.** *Let $G : X \rightrightarrows \mathbb{R}^d$ be outer semicontinuous and compact-convex valued. In addition to the conditions on W above, assume that for each x there exists $v \in G(x)$ such that $V'(x; v) + W(x; v) \leq 0$. Then there exists a trajectory of the differential inclusion $\dot{x} \in G(x)$ such that*

$$V(x(T)) - V(x(0)) + \int_0^T W(x(t), \dot{x}(t))dt \leq 0.$$

3.2 Functional Convergence of the Iteration Path

With our preliminaries out of the way, in this section we establish a general functional convergence theorem (Theorem 2) that applies to stochastic approximation-like algorithms that asymptotically approximate differential inclusions. By showing that we can represent both algorithms (5) and (6) in the stochastic approximation form our theorem requires, we then conclude that both schemes converge to the appropriate differential inclusion (Sec. 3.2.2).

3.2.1 A General Functional Convergence Theorem

Let $\{g_k\}_{k \in \mathbb{N}}$ be a collection of set-valued mappings $g_k : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$, and let $\{\alpha_k\}_{k \in \mathbb{N}}$ be a sequence of positive stepsizes. Now let $\{\xi_k\}_{k=1}^{\infty}$ be an arbitrary $\mathbb{R}^d$-valued sequence (the noise sequence), and consider the following iteration, which begins from the initial value $x_0 \in \mathbb{R}^d$:

$$x_{k+1} = x_k + \alpha_k[y_k + \xi_{k+1}], \quad \text{where } y_k \in g_k(x_k) \quad \text{for } k \geq 0. \tag{11}$$

In the coming subsection we show how this iteration encompasses both of our iteration schemes (5) and (6). For notational convenience, define the "times" $t_m = \sum_{k=1}^{m} \alpha_k$ as the partial stepsize sums, and let $x(\cdot)$ be the linear interpolation of the iterates x_k, that is,

$$x(t) := x_n + \frac{t - t_k}{t_{k+1} - t_k}(x_{k+1} - x_k) \quad \text{and} \quad y(t) = y_k \quad \text{for} \quad t \in [t_k, t_{k+1}). \tag{12}$$

Clearly this path satisfies $\dot{x}(t) = y(t)$ for almost all t and it is absolutely continuous on any compact interval. For $t \in \mathbb{R}_+$, define the time-shifted process $x^t(\cdot) = x(t + \cdot)$. Then we have the following general convergence theorem for the interpolated process (12) based on the iteration (11), where we recall that we metrize $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ with $\mathsf{d}(f, g) = \sum_{t=1}^{\infty} 2^{-t} \sup_{\tau \in [0,t]} \|f(\tau) - g(\tau)\| \wedge 1$.

Theorem 2. *Let the following conditions hold:*

(i) The iterates are bounded, i.e. $\sup_k \|x_k\| < \infty$ *and* $\sup_k \|y_k\| < \infty$.

(ii) The stepsizes are square summable but non-summable: $\sum_{k=1}^{\infty} \alpha_k = \infty$ *and* $\sum_{k=1}^{\infty} \alpha_k^2 < \infty$.

(iii) The weighted noise sequence is convergent: $\sum_{k=1}^{n} \alpha_k \xi_k \to v$ *for some* $v \in \mathbb{R}^d$ *as* $n \to \infty$.

(iv) There exists a closed-valued $H : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ *such that for all* $\{x_k\} \subset \mathbb{R}^d$ *satisfying* $\lim_k x_k = x$ *and all increasing subsequences* $\{n_k\}_{k \in \mathbb{N}} \subset \mathbb{N}$*, we have*

$$\lim_{n \to \infty} \operatorname{dist}\left(\frac{1}{n} \sum_{k=1}^{n} g_{n_k}(x_k), H(x) \right) = 0.$$

Then for any sequence $\{\tau_k\}_{k=1}^{\infty} \subset \mathbb{R}_+$*, the sequence of functions* $\{x^{\tau_k}(\cdot)\}$ *is relatively compact in* $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$*. If in addition* $\tau_k \to \infty$ *as* $k \to \infty$*, all limit points of* $\{x^{\tau_k}(\cdot)\}$ *in* $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ *satisfy*

$$\bar{x}(t) = \bar{x}(0) + \int_0^t y(\tau) d\tau \quad \text{for all } t \in \mathbb{R}_+$$

for a function $y : \mathbb{R}_+ \to \mathbb{R}^d$ *satisfying* $y(t) \in H(x(t))$.

The theorem is a generalization of Theorem 5.2 of Borkar [4], and the proof techniques are fairly similar. Consequently and for completeness, we provide its proof in Appendix A.

3.2.2 Limiting differential inclusion for stochastic prox-linear and gradient methods

With Theorem 2 in place, it is now of interest to show that both of the stochastic approximation schemes (5) and (6) can be represented by the general stochastic approximation scheme (11). As a consequence, we wish to verify that the stochastic prox-linear iteration (5) and the SGD iteration (6) satisfy the four conditions of Theorem 2. With this in mind, we introduce a bit of new notation before proceeding with our analysis. In analogy to the standard gradient mapping from both convex and composite optimization [25, 13], we define a stochastic gradient mapping G and consider its limits. In the stochastic proximal case, for fixed x we define

$$x_\alpha^+(s) := \operatorname*{argmin}_{y \in X} \left\{ f_x(y; s) + \varphi(y) + \frac{1}{2\alpha} \|y - x\|^2 \right\} \quad \text{and} \quad \mathsf{G}_\alpha(x; s) := \frac{1}{\alpha}(x - x_\alpha^+(s)), \tag{13a}$$

while for the subgradient case (6) we define

$$x_\alpha^+(s) := \operatorname*{argmin}_{y \in X} \left\{ \langle \mathsf{g}(x; s), y \rangle + \varphi(y) + \frac{1}{2\alpha} \|x - y\|^2 \right\} \quad \text{and} \quad \mathsf{G}_\alpha(x; s) := \frac{1}{\alpha}(x - x_\alpha^+(s)). \tag{13b}$$

To see that these updates are well-behaved (they are measurable in s [28, Lemma 1]), we present two lemmas on the subgradients of f and boundedness properties of G.

Lemma 3.6. *Let $f(x;s) = h(c(x;s);s)$ and $f(x) = \mathbb{E}_P[f(x;S)]$, where h and c satisfy Assumption A. Then*

$$\partial f(x;s) = \nabla c(x;s)\partial h(c(x;s)) \quad \text{and} \quad \partial f(x) = \mathbb{E}_P[\nabla c(x;S)\partial h(c(x;S);S)],$$

and $\partial f(\cdot) : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ is closed compact convex-valued and outer semicontinuous.

As the proof of Lemma 3.6 is somewhat technical and its results are not the main focus of this paper, we defer it to Appendix B.2.

Lemma 3.6 shows that $\partial f(x;s)$ is compact-valued and o.s.c., and we thus define the shorthand notation for the subgradients of $f + \varphi$ as

$$G(x;s) := \partial f(x;s) + \partial\varphi(x) \quad \text{and} \quad G(x) := \mathbb{E}_P[G(x;S)] = \int_{\mathcal{S}} \partial f(x;s) dP(s) + \partial\varphi(x), \tag{14}$$

both of which are outer-semicontinuous in x and compact-convex valued because φ is convex. Now we may show the boundedness properties of the gradient mappings (13).

Lemma 3.7. *For either of the updates* (13)*, we have $\|\mathsf{G}_\alpha(x;s)\| \le \|G(x;s)\|$.*

Proof For shorthand, write $x^+ = x_\alpha^+(x;s)$ and let $g = \mathsf{g}(x;s)$. By the definition of the optimality conditions for x^+, there exists a vector g^+ that, in the case of the update (13a), satisfies $g^+ \in \nabla c(x;s)\partial h(c(x;s) + \nabla c(x;s)^T(x^+ - x))$, and in the case of the update (13b), satisfies $g^+ = g$, and another vector $v^+ \in \partial\varphi(x^+)$ such that

$$\left\langle g^+ + \frac{1}{\alpha}(x^+ - x) + v^+, y - x^+ \right\rangle \ge 0 \quad \text{for all } y \in X.$$

Rearranging, we substitute $y = x$ to obtain

$$\left\langle g^+, x^+ - x \right\rangle + \frac{1}{\alpha}\left\|x - x^+\right\|^2 + \left\langle v^+, x^+ - x \right\rangle \le 0.$$

Using the monotonicity results that $\langle v^+, x - x^+\rangle \ge \langle \partial\varphi(x), x - x^+\rangle$ and $\langle g^+, x - x^+\rangle \ge \langle g, x - x^+\rangle$ because the subgradient mapping is monotone [19] for the functions $\varphi(x)$ and $y \mapsto f_x(y;s)$, we have

$$\left\langle g, x^+ - x \right\rangle + \frac{1}{\alpha}\left\|x - x^+\right\|^2 + \left\langle v, x^+ - x \right\rangle \le 0$$

for all $v \in \partial\varphi(x)$. The Cauchy-Schwartz implies $\|g + v\|\,\|x^+ - x\| \ge \frac{1}{\alpha}\|x - x^+\|^2$, which implies our desired result. □

By Lemma 3.7, in either of the updates (13), the vector $x_\alpha^+(s)$ is well-defined (even continuous in x and measurable in s). In order to define the population counterpart of the gradient mapping G_α, we require one more small result, which shows that the gradient mapping is locally bounded and integrable. To that end, for $x \in X$ and $\epsilon > 0$ we define the Lipschitz constants

$$L_\epsilon(x;s) := \sup_{x' \in X, \|x' - x\| \le \epsilon} \left\|G(x';s)\right\| \quad \text{and} \quad L_\epsilon(x) := \mathbb{E}_P[L_\epsilon(x;S)^2]^{\frac{1}{2}}.$$

These are well-behaved, as the following technical lemma shows (see Appendix B.3 for a proof).

Lemma 3.8. *Let Assumption A hold and $\epsilon \le 1$. Then $x \mapsto L_\epsilon(x;s)$ and $x \mapsto L_\epsilon(x)$ are upper semicontinuous on X and $L_\epsilon(x) < \infty$ for all $x \in X$.*

As a consequence of this lemma, we may define the mean gradient mapping

$$\overline{\mathsf{G}}_\alpha(x) := \mathbb{E}_P\left[\mathsf{G}_\alpha(x;S)\right] = \int_{\mathcal{S}} \mathsf{G}_\alpha(x;s)dP(s).$$

Moreover, it is now immediate that both the stochastic prox-linear (5) and projected stochastic subgradient algorithms (6) have the representation

$$x_{k+1} = x_k - \alpha_k \mathsf{G}_{\alpha_k}(x_k;S_k) = x_k - \alpha_k \overline{\mathsf{G}}_{\alpha_k}(x_k) - \alpha_k \xi_{\alpha_k}(x_k;S_k), \tag{15}$$

where the noise vector ξ has definition $\xi_\alpha(x;s) := \mathsf{G}_\alpha(x;s) - \overline{\mathsf{G}}_\alpha(x;s)$. By defining the filtration of σ-fields $\mathcal{F}_k$ by $\mathcal{F}_k := \sigma(x_0, S_1, \ldots, S_{k-1})$, we immediately have $x_k \in \mathcal{F}_k$ and that the noise sequence ξ is a square-integrable martingale difference sequence adapted to $\mathcal{F}_k$. Indeed, for any α and $\epsilon > 0$ we have

$$\|\mathsf{G}_\alpha(x;s)\| \le L_\epsilon(x;s) \;\text{ and }\; \left\|\overline{\mathsf{G}}_\alpha(x)\right\| \le L_\epsilon(x) \tag{16}$$

by Lemma 3.7 and the definition of the Lipschitz constant, and for any x and $\alpha > 0$ we have

$$\mathbb{E}_P\left[\|\xi_\alpha(x;S)\|^2\right] \le \mathbb{E}_P\left[\|\mathsf{G}_\alpha(x;S)\|^2\right] \le \mathbb{E}\left[L_\epsilon^2(x;S)\right] = L_\epsilon(x)^2, \tag{17}$$

because $\mathbb{E}[\mathsf{G}_\alpha] = \overline{\mathsf{G}}_\alpha$. In the context of our iterative procedures, for any $\alpha > 0$ we have

$$\mathbb{E}[\xi_\alpha(x_k;S_k) \mid \mathcal{F}_k] = 0 \;\text{ and }\; \mathbb{E}[\|\xi_\alpha(x_k;S_k)\|^2 \mid \mathcal{F}_k] \le L_\epsilon(x_k)^2.$$

In particular, the update form (15) shows that both the stochastic prox-linear iteration (5) and projected SGD (6) have the form (11) necessary for application of Theorem 2.

Functional convergence for the stochastic updates Now that we have the representation (15), it remains to verify that the mean gradient mapping $\overline{\mathsf{G}}$ and errors ξ satisfy the conditions necessary for application of Theorem 2. That is, we verify (i) bounded iterates, (ii) non-summable but square-summable stepsizes, (iii) convergence of the weighted error sequence, and (iv) the distance condition in the theorem. Condition (ii) is trivial (see Eq. (7)), so we ignore it. To address condition (i), we temporarily make the following assumption, noting that certainly the compactness of X is sufficient to it to hold.

Assumption C. *With probability 1, the iterates of the update schemes* (5) *and* (6) *are bounded,*

$$\sup_k \|x_k\| < \infty.$$

A number of conditions, such as almost supermartingale convergence guarantees explored by Robbins and Siegmund [27], are sufficient to demonstrate Assumption C holds. In particular, whenever Assumption C holds, we have that

$$\sup_k \sup_{\alpha>0} \left\|\overline{\mathsf{G}}_\alpha(x_k)\right\| \le \sup_k L_\epsilon(x_k) < \infty,$$

by Lemma 3.8 and inequality (16), because the supremum of an upper semicontinuous function on a compact set is finite. That is, condition (i) of Theorem 2 on the boundedness of x_k and y_k holds.

The error sequences ξ_{α_k} are also well-behaved for either the stochastic prox-linear updates (5) or the SGD updates (6). That is, condition (iii) of Theorem 2 is satisfied:

Lemma 3.9. *Let Assumptions A and C hold. Then with probability 1,* $\lim_n \sum_{k=1}^n \alpha_k \xi_{\alpha_k}(x_k; S_k)$ *exists and is finite.*

Proof Ignoring probability zero events, by Assumption C there is a (potentially random) constant $C < \infty$ such that $\|x_k\| \le C$ for all $k \in \mathbb{N}$. As $L_\epsilon(\cdot)$ is upper semicontinuous (Lemma 3.8), we know that $\sup\{L_\epsilon(x) \mid \|x\| \le C, x \in X\} < \infty$. Hence, using inequality (17), we have

$$\sum_{k=1}^\infty \mathbb{E}\left[\alpha_k^2 \left\|\xi_{\alpha_k}(x_k; S_k)\right\|^2 \mid \mathcal{F}_k\right] \le \sum_{k=1}^\infty \alpha_k^2 \sup_{\|x\| \le C, x \in X} L_\epsilon(x)^2 < \infty.$$

Standard convergence results for ℓ_2-summable martingale difference sequences (e.g. [10, Theorem 5.3.33]) immediately give the result. □

Finally, we verify the fourth technical condition Theorem 2 requires by constructing an appropriate closed-valued mapping $H : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$, which is identical for both algorithms (5) and (6). Recall the definition (14) of the outer semicontinuous mapping $G(x) = \mathbb{E}_P[\partial f(x; S)] + \partial\varphi(x)$. We then have the following limiting inclusion.

Lemma 3.10. *Let the sequence* $x_k \in X$ *satisfy* $x_k \to x \in X$. *Let* $\{n_k\} \subset \mathbb{N}$ *be an increasing sequence. Then, for either of the updates* (13),

$$\lim_{n\to\infty} \operatorname{dist}\left(\frac{1}{n}\sum_{k=1}^n \overline{\mathsf{G}}_{\alpha_{n_k}}(x_k), G(x) + \mathcal{N}_X(x)\right) = 0.$$

Proof Let $x_k^+(s)$ be shorthand for the result of the prox-linear (13a) or stochastic subgradient update (13b) when applied with the stepsize $\alpha = \alpha_{n_k}$. Then for any $\epsilon \in (0, 1)$, Lemma 3.7 shows that $\left\|x_k^+(s) - x_k\right\| \le \alpha_{n_k} L_\epsilon(x; s)$. By the standard (convex) optimality conditions for $x_k^+(s)$, there exists a vector $\mathsf{g}^+(x_k; s)$ such that, in the case of the update (13a), satisfies

$$\mathsf{g}^+(x_k; s) \in \nabla c(x_k; s)\partial h(c(x_k; s) + \nabla c(x_k; s)^T(x_k^+(s) - x); s),$$

and in the case of the update (13b), satisfies

$$\mathsf{g}^+(x_k; s) = \mathsf{g}(x_k; s) \in \nabla c(x_k; s)\partial h(c(x_k; s); s),$$

such that

$$\mathsf{G}_{\alpha_{n_k}}(x_k; s) \in \mathsf{g}^+(x_k; s) + \partial\varphi(x_k^+(s)) + \mathcal{N}_X(x_k^+(s)).$$

Let $v_k^+(s) \in \partial\varphi(x_k^+(s))$ and $w_k^+(s) \in \mathcal{N}_X(x_k^+(s))$ be the vectors such that

$$\mathsf{G}_{\alpha_{n_k}}(x_k; s) = \mathsf{g}(x_k; s) + v_k^+(s) + w_k^+(s).$$

The three set-valued mappings $x \mapsto \partial f(x; s)$, $x \mapsto \partial\varphi(x)$, and $x \mapsto \mathcal{N}_X(x)$ are outer semicontinuous (see Lemmas 3.1, 3.2, and 3.6). Since $x_k^+(s) \to x$ tends to x as $k \to \infty$, this outer semicontinuity thus implies

$$\operatorname{dist}\left(\mathsf{g}^+(x_k; s), \partial f(x; s)\right) \to 0, \quad \operatorname{dist}\left(v_k^+(s), \partial\varphi(x)\right) \to 0, \quad \text{and} \quad \operatorname{dist}\left(w_k^+(s), \mathcal{N}_X(x)\right) \to 0 \tag{18}$$

as $k \to \infty$. Because $x_k \to x$ and the Lipschitz constants $L_\epsilon(\cdot; s)$ are upper semicontinuous, Lemma 3.7 also implies that

$$\limsup_k \left\|\mathsf{g}^+(x_k; s) + v_k^+(s)\right\| \le L_\epsilon(x; s) \quad \text{and} \quad \limsup_k \left\|\mathsf{G}_{\alpha_{n_k}}(x_k; s)\right\| \le L_\epsilon(x; s).$$

By the triangle inequality, we thus obtain $\limsup_k \left\|w_k^+(s)\right\| \le 2L_\epsilon(x;s)$, and hence,

$$\operatorname{dist}\left(w_k^+(s), \mathcal{N}_X(x) \cap 2L_\epsilon(x;s) \cdot \mathbb{B}\right) \to 0.$$

The definition of the set-valued integral and $L_\epsilon(x) = \mathbb{E}[L_\epsilon(x;S)^2]^{\frac{1}{2}}$ yields that $\mathcal{N}_X(x) \cap 2L_\epsilon(x)\mathbb{B} \supset \int (\mathcal{N}_X(x) \cap 2L_\epsilon(x;s) \cdot \mathbb{B}) dP(s)$, and the definition of the set-valued integral and convexity of $\operatorname{dist}(\cdot,\cdot)$ (see Lemma B.1 in Appendix B.1 for rigorous justification of this step) imply that for any n

$$\begin{aligned}
&\operatorname{dist}\left(\frac{1}{n}\sum_{k=1}^n \overline{\mathsf{G}}_{\alpha_{n_k}}(x_k), G(x) + \mathcal{N}_X(x) \cap 2L_\epsilon(x) \cdot \mathbb{B}\right) \\
&\le \frac{1}{n}\sum_{k=1}^n \int \operatorname{dist}\left(\mathsf{G}_{\alpha_{n_k}}(x_k;s), \partial f(x;s) + \partial\varphi(x) + \mathcal{N}_X(x) \bigcap 2L_\epsilon(x;s) \cdot \mathbb{B}\right) dP(s). \qquad (19)
\end{aligned}$$

We now bound the preceding integral. By the definition of Minkowski addition and the triangle inequality, we have the pointwise convergence

$$\begin{aligned}
&\operatorname{dist}\left(\mathsf{G}_{\alpha_{n_k}}(x_k;s), \partial f(x;s) + \partial\varphi(x) + \mathcal{N}_X(x) \cap 2L_\epsilon(x;s) \cdot \mathbb{B}\right) \\
&\le \operatorname{dist}\left(\mathsf{g}(x_k;s), \partial f(x;s)\right) + \operatorname{dist}\left(v_k^+(s), \partial\varphi(x)\right) + \operatorname{dist}\left(w_k^+(s), \mathcal{N}_X(x) \cap 2L_\epsilon(x;s) \cdot \mathbb{B}\right) \to 0
\end{aligned}$$

as $k \to \infty$ by the earlier outer semicontinuity convergence guarantee (18). For suitably large k, each of the terms in the preceding sum is upper bounded by $2L_\epsilon(x;s)$, which is square integrable by Lemma 3.8. Lebesgue's dominated convergence theorem thus implies that the individual summands in expression (19) converge to zero as $k \to \infty$, and the simple analytic fact that the Cesáro mean $\frac{1}{n}\sum_{k=1}^n a_k \to 0$ if $a_k \to 0$ as $k \to \infty$ gives the result. □

With this lemma, we may now show the functional convergence of the stochastic linear prox (5) and stochastic gradient (6) update schemes. We have verified that each of the conditions (i)–(iv) of Theorem 2 hold with the mapping $H(x) = -\mathcal{N}_X(x) - G(x)$. Indeed, H is closed-valued and outer-semicontinuous as $G(\cdot)$ is convex compact o.s.c. and $\mathcal{N}_X(\cdot)$ is closed and o.s.c. Thus, with slight abuse of notation, let $x(\cdot)$ be the linear interpolation (12) of the iterates x_k for either the stochastic prox-linear algorithm or the stochastic subgradient algorithm, where we recall that $x^t(\cdot) = x(t + \cdot)$. Then we have the following convergence theorem.

Theorem 3. *Let Assumptions A and C hold. With probability one over the random sequence $S_1, S_2, \ldots$, we have the following. For any sequence $\{\tau_k\}_{k=1}^\infty$, the function sequence $\{x^{\tau_k}(\cdot)\}$ is relatively compact in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$. In addition, for any sequence $\tau_k \to \infty$ with $\tau_k \to \infty$, any limit point of $\{x^{\tau_k}(\cdot)\}$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ satisfies the integral equation*

$$\bar{x}(t) = \bar{x}(0) + \int_0^t y(\tau) d\tau \quad \text{for all } t \in \mathbb{R}_+, \quad \text{where} \quad y(\tau) \in -G(x(\tau)) - \mathcal{N}_X(x(\tau)).$$

3.3 Properties of the Limiting Differential Inclusion

Theorem 3 establishes that both the stochastic prox-linear (5) and subgradient (6) procedures have sample paths asymptotically approximated by the differential inclusion

$$\dot{x} \in -G(x) - \mathcal{N}_X(x) \quad \text{where} \quad G(x) = \partial f(x) + \partial\varphi(x)$$

for the objective $f(x) = \mathbb{E}[h(c(x;S);S)]$. To establish convergence of the iterates x_k themselves, we must understand the limiting properties of trajectories of the preceding differential inclusion. As we

see presently, the structure of G is amenable to analysis, the differential inclusion (9) has a unique solution from any $x_0 \in X$, and it admits a reasonably simple Lyapunov convergence inequality.

The first step in this is the following lemma, which shows that the composite function $f + \varphi$ is not too non-convex (in the parlance of Rockafellar and Wets [29], it lower $\mathcal{C}^2$), which in turn allows us to demonstrate uniqueness of solutions to $\dot{x} \in -G(x) - \mathcal{N}_X(x)$.

Lemma 3.11. *Let $K \subset X$ be compact convex and satisfy $\sup_{x \in K} \|x\| \leq B$. Then for any x_0 and $\lambda \geq \Lambda(B) := \mathbb{E}[\gamma(B;S)\beta(B;S)]$, the function $f_\lambda(x) := f(x) + \frac{\lambda}{2}\|x - x_0\|^2$ is convex on K.*

Proof The proof mimics Proposition 2.1 of Drusvyatskiy and Kempton [11]. Fix $s \in \mathcal{S}$. Then by the definition of $\Lambda(B)$ and $\Lambda(B;s) = \gamma(B;s)\beta(B;s)$, we have that $h(\cdot;s)$ is $\gamma(B;s)$-Lipschitz at least on the ball of radius $\sup\{\|c(x;s)\| : x \in X, \|x\| \leq B\}$. Thus $h^*(y;s) = \sup_w\{w^T y - h(w;s)\}$ has domain contained in $\{y : \|y\| \leq \gamma(B;s)\}$. Noting that for any $w \in \operatorname{dom} h^*(\cdot;s)$, the function $x \mapsto \langle w, c(x;s)\rangle$ is thus $\Lambda(B;s)$-smooth, we obtain

$$\langle w, c(x;s)\rangle + \frac{\Lambda(B;s)}{2}\|x - x_0\|^2$$

is convex in x for any x_0, and

$$h(c(x;s);s) + \frac{\gamma(B;s)\beta(B;s)}{2}\|x - x_0\|^2 = \sup_w \left\{ c(x;s)^T w - h^*(w;s) + \frac{\gamma(B;s)\beta(B;s)}{2}\|x - x_0\|^2 \right\}.$$

As the supremum of convex functions, the left term is thus convex. Set $\lambda = \mathbb{E}[\gamma(B;S)\beta(B;S)]$. □

In combination with the uniqueness guarantee of Lemma 3.4, this lemma is the key result that allows us to prove the theorem to come on the differential inclusion (9). Recall that a function f is coercive if $f(x) \to \infty$ as $\|x\| \to \infty$. Then if we define the minimal subgradient

$$\mathsf{g}^\star(x) := \operatorname*{argmin}_g \left\{ \|g\|^2 \mid g \in \partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x) \right\} = \pi_{G(x)+\mathcal{N}_X(x)}(0),$$

we obtain the following convergence theorem on the differential inclusion.

Theorem 4. *Assume that $f + \varphi + \mathbb{I}_X$ is coercive. Let $x(\cdot)$ be a solution to the differential inclusion $\dot{x} \in -\partial f(x) - \partial\varphi(x) - \mathcal{N}_X(x)$ initialized at $x(0) \in X$. Then $x(t)$ exists for all times $t \in \mathbb{R}_+$, $\sup_t \|x(t)\| < \infty$, $x(t)$ is Lipschitz continuous in t, $x(t) \in X$, and*

$$f(x(t)) + \varphi(x(t)) + \int_0^t \|\mathsf{g}^\star(x(\tau))\|^2 \, d\tau \leq f(x(0)) + \varphi(x(0)).$$

We prove the theorem in Section 3.3.1 to come, giving a few corollaries to show that solutions to the differential inclusion converge to stationary points of $f + \varphi$. We first have

Corollary 3.1. *Let $x(\cdot)$ be a solution to the differential inclusion $\dot{x} \in -G(x) - \mathcal{N}_X(x)$ and assume that for some $t > 0$ we have $f(x(t)) = f(x(0))$. Then $\mathsf{g}^\star(x(\tau)) = 0$ for all $\tau \in [0,t]$.*

Proof By Theorem 4, we have that $\int_0^t \|\mathsf{g}^\star(x(\tau))\|^2 \, d\tau = 0$, so that $\mathsf{g}^\star(x(\tau)) = 0$ for almost every $\tau \in [0,t]$. The continuity of $x(\cdot)$ and outer semi-continuity of G extend this to all τ. □

In addition, we can show that all cluster points of any trajectory solving the differential inclusion (9) are stationary. First, we recall the following definition.

Definition 3.1. *Let $\{x(t)\}_{t\geq 0}$ be a trajectory. A point x_∞ is a* cluster point *of $x(t)$ if there exists an increasing sequence $t_n \to \infty$ such that $x(t_n) \to x_\infty$. Let $\mathcal{T}_\epsilon(x_\infty) = \{t \in \mathbb{R}_+ \mid \|x(t) - x_\infty\| \leq \epsilon\}$. Let μ be Lebesgue measure on $\mathbb{R}_+$. A point x_∞ is an* almost cluster point *of $x(\cdot)$ if*

$$\mu(\mathcal{T}_\epsilon(x_\infty) \cap [T, \infty)) = \infty \quad \text{for all } \epsilon > 0 \text{ and } T < \infty.$$

It is immediate that all almost cluster points are also cluster points. Theorem 4 implies that cluster points of solutions to $\dot{x} \in -G(x) - \mathcal{N}_X(x)$ are also almost cluster points, because the trajectory $x(\cdot)$ is Lipschitz (cf. [1, Proposition 6.5.1]). We also have the following observation.

Corollary 3.2. *Let x_∞ be a cluster point of the trajectory $x(\cdot)$ for $\dot{x} \in -G(x) - \mathcal{N}_X(x)$. Then x_∞ is stationary, meaning that $\mathsf{g}^\star(x_\infty) = 0$.*

Proof By the remark preceding the statement of the corollary, x_∞ is also an almost cluster point of the trajectory. Let ϵ_n, δ_n be sequences of positive numbers converging to 0. Because $f(x(t)) + \varphi(x(t))$ converges to $f(x_\infty) + \varphi(x_\infty)$ (because the sequence is decreasing and $f + \varphi$ is continuous), we have $\int \|\mathsf{g}^\star(x(t))\|^2 \, dt < \infty$. Moreover, there exist increasing T_n such that

$$\int_{\mathcal{T}_{\epsilon_n}(x_\infty) \cap [T_n, \infty)} \|\mathsf{g}^\star(x(t))\|^2 \, dt \leq \delta_n.$$

Because $\mu(\mathcal{T}_{\epsilon_n}(x_\infty) \cap [T_n, \infty)) = \infty$, there must exist an increasing sequence $t_n \geq T_n$, $t_n \in \mathcal{T}_{\epsilon_n}(x_\infty)$, such that $\|\mathsf{g}^\star(x(t_n))\|^2 \leq \delta_n$. By construction $x(t_n) \to x_\infty$, and we have a subsequence $\mathsf{g}^\star(x(t_n)) \to 0$. The outer semi-continuity of $x \mapsto G(x) + \mathcal{N}_X(x)$ implies that $0 \in G(x_\infty) + \mathcal{N}_X(x_\infty)$. □

3.3.1 Proof of Theorem 4

Our argument proceeds in three main steps. For shorthand, we define $F(x) = f(x) + \varphi(x)$. Our first step shows that the function $V(x) := F(x) + \mathbb{I}_X(x) - \inf_{y \in X} F(y)$ is a Lyapunov function for the differential inclusion (9), where we take the function W in Lemma 3.5 to be $W(x, v) = \|v\|^2$. Once we have this, then we can use the existence result of Lemma 3.3 to show that a solution $x(\cdot)$ exists in a neighborhood of 0. The uniqueness of trajectories (Lemma 3.4) then implies that the trajectory x is non-increasing for V, which then—combined with the assumption of coercivity of $F + \mathbb{I}_X$—implies that the trajectory x is bounded and we can extend uniquely it to all of $\mathbb{R}_+$.

Part 1: A Lyapunov function To develop a Lyapunov function, we compute (approximate) directional derivatives of $f + \varphi$. Recalling that $f_x(y) := \int h(c(x; s) + \nabla c(x; s)^T (y - x); s) dP(s)$ gives the following approximation result, an immediate consequence of the Lipschitzian guarantees of Assumption A.

Lemma 3.12. *Let f_x be as above and $B > \|x\|$. Then for all $y \in X$ with $\|x - y\| \leq 1$ and $\|y\| \leq B$,*

$$|f(y) - f_x(y)| \leq \frac{\mathbb{E}[\gamma(B; S)\beta(B; S)]}{2} \|x - y\|^2 .$$

We also have the following essentially standard result on directional derivatives of convex functions.

Lemma 3.13 (Hiriart-Urruty and Lemaréchal [19], Chapter VI.1)**.** *Let h be convex and $g^\star = \pi_{\partial h(x)}(0) = \operatorname{argmin}_{g \in \partial h(x)}\{\|g\|\}$. Then the directional derivative satisfies $h'(x; -g^\star) = -\|g^\star\|^2$.*

Now, take $\mathsf{g}^\star(x)$ as in the statement of the theorem and $V(x) = f(x) + \varphi(x) + \mathbb{I}_X(x) - \inf_{y\in X}\{f(y) + \varphi(y)\}$; we claim that

$$V'(x; -\mathsf{g}^\star(x)) \le - \left\|\mathsf{g}^\star(x)\right\|^2. \tag{20}$$

Before proving the claim (20), we note that the condition (20) is identical to that in Lemma 3.5 on monotone trajectories of differential inclusions. Thus we obtain that there exists a solution $x(\cdot)$ to the differential inclusion $\dot{x} \in -G(x) - \mathcal{N}_X(x)$ defined on $[0, T]$ for some $T > 0$, where $x(\cdot)$ satisfies

$$f(x(t)) + \varphi(x(t)) + \mathbb{I}_X(x(t)) \le f(x(0)) + \varphi(x(0)) - \int_0^t \left\|\mathsf{g}^\star(x(\tau))\right\|^2 d\tau \tag{21}$$

for all $t \in [0, T]$. We return now to prove the claim (20). Let $B \ge \|x\|$, $v \in \mathbb{R}^d$, and $t < 1/\|v\|$ be otherwise arbitrary, so we have

$$|f(x+tv)+\varphi(x+tv)-f(x)-\varphi(x)| \le |f_x(x+tv)+\varphi(x+tv)-f(x)-\varphi(x)| + \frac{t^2\mathbb{E}[\beta(B;S)\gamma(B;S)]}{2}\|v\|^2$$

by Lemma 3.12. Because φ is convex and the error in the approximation f_x of f is second-order, taking limits as $u \to v, t \to 0$, we have for any fixed $x \in X$ that

$$\begin{aligned}
&\liminf_{t\downarrow 0, u\to v} \frac{F(x+tu) + \mathbb{I}_X(x+tu) - F(x)}{t} \\
&= \liminf_{t\downarrow 0} \frac{f_x(x+tv) + \varphi(x+tv) + \mathbb{I}_X(x+tv) - f_x(x) - \varphi(x)}{t} = \sup_{g \in \partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x)} \langle g, v\rangle,
\end{aligned}$$

where we have used that the subgradient set of $y \mapsto f_x(y)$ at $y = x$ is $\partial f(x)$ and the definition of the normal cone to X at x. Applying Lemma 3.13 with $v = -\mathsf{g}^\star(x)$ gives claim (20).

Part 2: Uniqueness of trajectories We now use Lemma 3.4 to show that solutions to $\dot{x} \in -G(x) - \mathcal{N}_X(x)$ have unique trajectories. Lemma 3.11 shows that for all $B < \infty$, the function $f(x) + \varphi(x) + \frac{\lambda}{2}\|x\|^2$ is convex on the set $X \cap \{x : \|x\| \le B\}$ for $\lambda \ge \Lambda(B)$. Thus for any points x_1, x_2 satisfying $\|x_i\| \le B$ and any $g_i \in \partial f(x_i) + \partial\varphi(x_i) + \mathcal{N}_X(x_i)$, we have

$$\langle g_1 + \lambda x_1 - g_2 - \lambda x_2, x_1 - x_2\rangle \ge 0$$

by Lemma 3.11 and that subgradients of convex functions are increasing [19, Ch. VI]. Rearranging, we have

$$\langle -g_1 + g_2, x_1 - x_2\rangle \le \lambda \|x_1 - x_2\|^2 \quad \text{for } g_i \in G(x_i) + \mathcal{N}_X(x_i).$$

This is equivalent to the condition of Lemma 3.4, so that for any B and any interval $[0, T]$ for which the trajectory $x(t)$ satisfies $\|x(t)\| \le B$ on $t \in [0, T]$, we have that the trajectory is unique. In particular, we have that the Lyapunov inequality (21) is satisfied on the interval over which the trajectory $\dot{x} \in -G(x) - \mathcal{N}_X(x)$ is defined.

Part 3: Extension to all times Lastly, we argue that we may take $T \to \infty$. For any fixed $T < \infty$, we know that $f(x(T)) + \varphi(x(t)) \le f(x(0)) + \varphi(x(0))$, and the coercivity of $f + \varphi$ over X implies that $\|x(t)\|$ must be uniformly bounded on this trajectory. Thus, there exists some $B < \infty$ such that $\|x(t)\| \le B$ for all $t \in [0, T]$. But then the compactness of $\partial f(x) + \partial\varphi(x)$ for $x \in X \cap \{y : \|y\| \le B\}$ implies that the projection of 0 onto the set $\partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x)$ has bounded norm (because $0 \in \mathcal{N}_X(x)$). Thus the condition on existence of paths for all times T in Lemma 3.3 applies, so that we may take $T \to \infty$.

The Lipschitz condition on $x(t)$ is an immediate consequence of the boundedness of the subgradient sets $\partial f(x) + \partial\varphi(x)$ for bounded x.

3.4 Almost Sure Convergence to Stationary Points

Thus far we have shown that the limit points of the stochastic iterations (5) and (6) are asymptotically equivalent to the differential inclusion (9) (Theorem 3) and that solutions to the differential inclusion have certain uniqueness and convergence properties (Theorem 4). Because the stochastic iterates x_k are in general never exactly on a solution path $x(t)$—there is always some noise—we provide an additional argument showing more subtle stability properties of the differential inclusion to perturbations, which is the purpose of this section. In particular, we show that all cluster points of the iterates x_k of either procedure (5) or (6) are stationary and that $f(x_k) + \varphi(x_k)$ converges.

To provide a starting point, we state the result, which is our main convergence theorem.

Theorem 5. *Let Assumptions A, B, and C hold. Assume that $f + \varphi + \mathbb{I}_X$ is coercive. Then with probability 1 all cluster points of the sequence $\{x_k\}_{k=1}^\infty$ belong to the stationary set $X^\star$ and $f(x_k) + \varphi(x_k)$ converges.*

The proof of the theorem relies on a stability analysis we perform in Section 3.4.1. In order to illustrate the theorem, however, we first give a few examples that show how to satisfy Assumption C, that is, that $\sup_k \|x_k\| < \infty$. We remark in passing that Theorem 1 is an immediate consequence of Theorem 5, because if X is compact then the iterates x_k are bounded and $f + \varphi + \mathbb{I}_X$ is coercive.

Conditions for boundedness of the iterates We may develop more subtle examples by considering the joint properties of the regularizer φ and objectives $f(x; S)$ in the stochastic updates of our methods. Rather than providing an exhaustive characterization of boundedness, we simply provide two examples for motivation of Assumption C, focusing for simplicity on the stochastic (Fréchet) subgradient update (6) in the unconstrained case when $X = \mathbb{R}^d$.

First, let us assume that $\varphi(x) = \frac{\lambda}{2}\|x\|^2$, simple ℓ_2- or Tikhonov regularization, common in numerous machine learning, statistics, and inverse problems. In addition, let us assume that $f(x; s) = h(c(x; s); s)$ is $L(s)$-Lipschitz in x, where $L := \mathbb{E}[L(S)^2]^{\frac{1}{2}} < \infty$, and so that $\|\mathsf{g}(x; s)\| \le L(s)$. This regularization is sufficient to guarantee boundedness with probability 1:

Lemma 3.14. *Let the conditions of the preceding paragraph hold. Assume that $\mathbb{E}[L(S)^2] < \infty$. Then with probability 1, $\sup_k \|x_k\| < \infty$.*

We provide the proof of Lemma 3.14 using a martingale argument in Appendix B.4.

As a second example, we show how more quickly growing regularization functions φ may also yield boundedness of iterates in an extension of this result. First, we recall the following

Definition 3.2. *A function $\varphi : \mathbb{R}^d \to \mathbb{R}$ is* β-coercive *if $\lim_{\|x\|\to\infty} \varphi(x)/\|x\|^\beta = \infty$.*

A standard result [19, Chapter IV.3] is that if a convex function φ is β-coercive on $\mathbb{R}^d$, then $\|v(x)\| / \|x\|^{\beta-1} \to \infty$ for $v(x) \in \partial\varphi(x)$ whenever $\|x\| \to \infty$. Intuitively, we expect that if $\varphi(x)$ grows quickly enough as $\|x\| \to \infty$, then the the iterates x_k also remain bounded. To make the coming argument simpler, we make the reasonable assumption that φ is regular enough that there exists a constant $\lambda \in (0, 1]$ such that $\varphi(x) \ge \varphi(\lambda x)$ for x with $\|x\|$ sufficiently large; $\varphi(x) = \|x\|^\beta$ satisfies this inequality with $\lambda = 1$.

Lemma 3.15. *Let φ be β-coercive and satisfy the above regularity property. Assume that for all $s \in \mathcal{S}$, $x \mapsto f(x; s) = h(c(x; s); s)$ is $L(1 + \|x\|^\nu)$-Lipschitz in a neighborhood of x, where $L < \infty$ is some constant, and $\nu < \beta - 1$. Then $\sup_k \|x_k\| < \infty$.*

See Appendix B.5 for a proof of the lemma.

Lemmas 3.14 and 3.15 give two concrete examples that are sufficient to guarantee boundedness of the iterates; this motivates our belief that generally, Assumption C is not too onerous.

3.4.1 Stability of the differential inclusion and proof of Theorem 5

We now provide the stability analysis necessary for the proof of Theorem 5. For $\rho \in \mathbb{R}$, let A_ρ denote the sublevel sets of our objective function $F(x) = f(x) + \varphi(x)$,

$$A_\rho := \{x \in X : F(x) \le \rho\}.$$

We denote the δ-neighborhood of a set A by $A^\delta := A + \delta\mathbb{B}$, and we make the following standard stability definition for trajectories of a differential inclusion $\dot{x} \in H(x)$.

Definition 3.3 (Stability)**.** *A set $A \subset X$ is* locally stable *if for all $\delta > 0$, there exists $\delta' > 0$ such that any trajectory of $\dot{x} \in H(x)$ with initial point $x(0) \in A^{\delta'}$ satisfies $x(t) \in A^\delta$ for all t. If $\limsup_t \operatorname{dist}(x(t), A) = 0$ for such trajectories, then A is* locally asymptotically stable.

One sometimes appends the names in Definition 3.3 with "in the sense of Lyapunov" [21, 4].

Roughly, our approach is to show that (most of) the sublevel sets A_ρ are locally asymptotically stable, which means that eventually the perturbations of the stochastic iteration are negligible, and the iterates $f(x_k) + \varphi(x_k)$ must converge with the differential inclusion. Our argument bears similarities to the standard ODE method as explicated by Kushner and Yin [21, Theorem 5.2.1], though non-smoothness of f forces somewhat more care in our case. Before continuing, recall the definition (8) of the set of stationary points $X^\star = \{x \in X : 0 \in \partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x)\}$ and that $\mathsf{g}^\star(x) = \operatorname{argmin}_g\{\|g\| : g \in \partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x)\}$.

The first step of our stability analysis shows that sublevel sets are locally asymptotically stable as long as some neighborhood of the set contains no stationary points.

Lemma 3.16. *Let $\rho \in F(X) = \{f(x) + \varphi(x) : x \in X\}$ and suppose that there exists some $\epsilon > 0$ such that $A_{\rho+\epsilon} \setminus A_\rho$ contains no stationary points. Then the sublevel set A_ρ is locally asymptotically stable for the differential inclusion* (9)*.*

Proof We claim that $A_{\rho+\epsilon}$ is a domain of attraction for A_ρ, meaning that any trajectory $x(\cdot)$ with $x(0) \in A_{\rho+\epsilon}$ satisfies $x(t) \to A_\rho$. For the sake of contradiction, consider a trajectory $x(\cdot)$ beginning from a point $x(0) \in A_{\rho+\epsilon}$ with $\liminf_t \operatorname{dist}(x(t), A_\rho) > 0$. (The monotone trajectory for the solution of the differential inclusion (9), by Theorem 4, shows that if $\liminf_t \operatorname{dist}(x(t), A_\rho) = 0$, then $x(t) \to A_\rho$.) In particular, as all cluster points of trajectories are stationary (Cor. 3.2), there exists a point x satisfying $\delta > \operatorname{dist}(x, A_{\rho'}) > 0$ such that $\mathsf{g}^\star(x) = 0$. This contradicts the assumptoin that $A_{\rho+\epsilon} \setminus A_\rho$ contains no stationary points. Finally, by compactness of the sublevel sets (because $f+\varphi+\mathbb{I}_X$ is coercive) and continuity of $F = f+\varphi$, there exists some $\delta > 0$ such that $A_\rho^\delta \subset A_{\rho+\epsilon}$. $\square$

We note the following consequence of Assumption B—that the image $F(X^\star) = \{f(x) + \varphi(x) : x \in X^\star\}$ is countable—which with Lemma 3.16 shows that for almost all $\rho \in \mathbb{R}$, the sublevel set A_ρ is locally asymptotically stable.

Lemma 3.17. *Let Assumption B hold. For any $\rho \in \mathbb{R}$ and any $\epsilon > 0$, there exist $\rho' \in [\rho, \rho+\epsilon]$ and $\rho'' > \rho'$ such that*

$$A_{\rho''} \setminus A_{\rho'}$$

contains no stationary points.

Proof Suppose for the sake of contradiction that this does not hold at some $\rho \in F(X)$. Then there is an $\epsilon > 0$ such that for all $\rho' \in [\rho, \rho+\epsilon]$ and all $\rho'' > \rho'$, the set $A_{\rho''} \setminus A_{\rho'}$ contains a stationary point. Now, we know by assumption that there is some $\rho' \in [\rho, \rho+\epsilon]$ such that the boundary

$\operatorname{bd} A_{\rho'} \subset F^{-1}(\{\rho'\})$ contains no stationary points, because the image $F(X^\star)$ is countable. Then if $A_{\rho''} \setminus A_{\rho'}$ contains a stationary point for all $\rho'' > \rho'$, we can construct a sequence of points x_n that are stationary, that is, $\mathsf{g}^\star(x_n) = 0$ while $x_n \notin A_{\rho'}$, but $x_n \to A_{\rho'}$. The outer semi-continuity (Lemma 3.6) of $x \mapsto \partial f(x) + \partial\varphi(x) + \mathcal{N}_X(x)$ then implies that $\operatorname{bd} A_{\rho'}$ contains a stationary point. This is a contradiction, which yields our claimed lemma. □

With the help of previous lemma, we can prove the main technical result in this section.

Lemma 3.18. *Let $x(\cdot)$ be the linear interpolation* (12) *of the stochastic prox-linear* (5) *or subgradient* (6) *updates. With probability 1, $F(x(t))$ converges as $t \to \infty$.*

Proof Fix ω, the sample path of the observations $S_1(\omega), S_2(\omega), \ldots$. Theorem 3 implies that the interpolated function $x(\cdot)$ converges to some path $\bar{x}$ satisfying the differential inclusion (9). For a set $A \subset X$, let $\mathcal{D}_A$ be the attractor for the set A, that is, those points $x_0 \in X$ such that the trajectory of the differential inclusion (9) initialized at $x(0) = x_0$ satisfy $\limsup_{t\to\infty} \operatorname{dist}(x(t), A) = 0$. We make the following claim. Let $\rho \in \mathbb{R}$, $\delta > 0$ be such that $\mathcal{D}_{A_\rho} \supset A_{\rho+\delta}$ and $A_{\rho+\delta} \setminus A_\rho$ contains no stationary points. Then if $x(\cdot)$ enters $\mathcal{D}_{A_\rho}$ infinitely often,

$$\limsup_{t\to\infty} \{x(t)\} \subset A_\rho. \tag{22}$$

Equivalently, all cluster points of the path $t \mapsto x(t)$ belong to A_ρ. Before proving the claim (22), let us show how it yields a quick proof of the lemma. Define $\rho^* := \liminf_{t\to\infty} F(x(t))$. Then claim (22) shows that all cluster points of $x(t)$ as $t \to \infty$ belong to A_{ρ^*}, and as any sequence $x(t_k)$ for $t_k \to \infty$ has a further subsequence that converges, we must have $F(x(t)) \to \rho^*$.

To see the claim (22), let $0 < \delta_2 < \delta_1 < \delta$. Note that, since $x(\cdot)$ enters $\mathcal{D}_{A_\rho}$ infinitely often, there must exist some $x_0 \in \mathcal{D}_{A_\rho}$ such that x_0 is a cluster point of $x(\cdot)$. By Theorem 3, there must thus exist a some sequence $\{h_k\}_{k=1}^\infty$, $h_k \to \infty$, such that the shifted interpolated process $x^{h_k}(\cdot)$ converges to a trajectory $\bar{x}$ of the inclusion (9) satisfying $\bar{x}(0) = x_0$. (Take the sequence of times h_k to be such that $x(h_k) \to x_0$.) By the definition 3.3 of local asymptotic stability and the attractor $\mathcal{D}_{A_\rho}$, we know that the $\bar{x}(\cdot)$ converges to the set A_ρ so that $\limsup_{t\to\infty} \operatorname{dist}(\bar{x}(t), A_\rho) = 0$. In particular, Theorem 3 yields the functional convergence

$$\lim_{k\to\infty} \sup_{t\in[0,T]} |x(h_k + t) - \bar{x}(t)| \to 0,$$

and thus $x(\cdot)$ enters the set $A_{\rho+\delta_2}$ infinitely often.

Now we show that the interpolation $x(\cdot)$ cannot exit the set $A_{\rho+\delta_1}$ infinitely often. Suppose to the contrary that $x(\cdot)$ exits $A_{\rho+\delta_1}$ infinitely often. Let $\{h'_k\}_{k=1}^\infty$ and $\{h_k\}_{k=1}^\infty$ to be two sequences satisfying $h'_k < h_k$, $\lim_{k\to\infty} h_k = \lim_{k\to\infty} h'_k = \infty$, and that

$$F(x(h'_k)) = \rho + \delta_2, \quad F(x(h_k)) = \rho + \delta_1, \quad \text{and} \quad \rho + \delta_2 < F(x) < \rho + \delta_1 \text{ for } \quad x \in (h'_k, h_k). \tag{23}$$

To see that sequences satisfying condition (23) exist, we take h_k and h'_k as the "traversals" of the interval $[\rho + \delta_2, \rho + \delta_1]$: since $x(\cdot)$ enters $A_{\rho+\delta_2}$ infinitely often and exits $A_{\rho+\delta_1}$ infinitely often, and $x(\cdot)$ and $F(\cdot)$ are continuous, we know that there exist increasing sequences $\tilde{h}_k$ and $\tilde{h}'_k$ such that

$$F(x(\tilde{h}'_k)) = \rho + \delta_2, \quad F(x(\tilde{h}_k)) = \rho + \delta_1 \quad \text{and} \quad \tilde{h}'_k < \tilde{h}_k.$$

Then we define the last entrance and first subsequent exit times

$$h'_k := \sup\{h \in [\tilde{h}'_k, \tilde{h}_k] : f(x(h)) \le \rho + \delta_2\} \quad \text{and} \quad h_k := \inf\{h \in [h'_k, \tilde{h}_k] : f(x(h)) \ge \rho + \delta_1\}.$$

The continuity of F and $x(\cdot)$ show that the conclusion (23) holds.

We now argue that any sequence h_k, h_k' satisfying condition (23) has $\liminf_{k\to\infty}(h_k' - h_k) > 0$. The continuity of $F(\cdot)$ immediately implies $\liminf_{k\to\infty} \|x(h_k') - x(h_k)\| > 0$. Now, note the following fact: for any $m \leq n \in \mathbb{N}$, we have

$$\|x(t_n) - x(t_m)\| = \left\| \sum_{i=m+1}^{n} \alpha_i \overline{\mathsf{G}}_{\alpha_i}(x_i) + \sum_{i=m+1}^{n} \alpha_i \xi_i \right\| \leq (t_n - t_m) \sup_i \left\| \overline{\mathsf{G}}_{\alpha_i}(x_i) \right\| + \left\| \sum_{i=m+1}^{n} \alpha_i \xi_i \right\|.$$

Let $M = \sup_i \left\| \overline{\mathsf{G}}_{\alpha_i}(x_i) \right\| < \infty$ (use Lemmas 3.7 and 3.8 to see that $M < \infty$). Now, for $h \in \mathbb{R}_+$ define the time indices

$$k_<(h) = \max\{k \in \mathbb{N} : t_k \leq h\} \quad \text{and} \quad k_>(h) = \min\{k \in \mathbb{N} : t_k \geq h\}.$$

Then the statement $\liminf_{k\to\infty}(h_k - h_k') > 0$ is equivalent to $\liminf_{k\to\infty}(t_{k_>(h_k)} - t_{k_<(h_k')}) > 0$. Note that Lemma 3.9 implies that $\lim_{m\to\infty} \sup_{n\geq m} \|\sum_{i=m+1}^{n} \alpha_i \xi_i\| = 0$, and thus we obtain that for any $\epsilon > 0$ and all suitably large m, we have

$$(t_n - t_m)M \geq \|x(t_n) - x(t_m)\| - \left\| \sum_{i=m+1}^{n} \alpha_i \xi_i \right\| \geq \|x(t_n) - x(t_m)\| - \epsilon.$$

Then for any $\epsilon > 0$, for all suitably large k we have the bound

$$\left\| x(h_k') - x(h_k) \right\| \leq \max\left\{ \|x(t_n) - x(t_m)\| : n, m \in [k_<(h_k'), k_>(h_k)] \right\} \leq \left(t_{k_>(h_k)} - t_{k_>(h_k')} \right) M + \epsilon.$$

This implies that $\liminf_{k\to\infty}(h_k - h_k') > 0$ as desired.

Let $\liminf_{k\to\infty}(h_k - h_k') = 2c > 0$ for a constant c. Then for all suitably large k we can find points $l_k, r_k \in [h_k', h_k]$ such that $r_k - l_k \in [c, 4c]$ and $F(x(r_k)) \geq F(x(l_k))$. By Theorem 3, the functional sequence $x^{l_k}(\cdot)$ has a subsequence converging to a trajectory of the inclusion (9) in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$, which we denote $\bar{x}'(\cdot)$. Without loss of generality (by considering appropriate subsequences as necessary), we assume the entire sequence $x^{l_k}(\cdot) \to \bar{x}'(\cdot)$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ with $\bar{x}'(0) = x_0'$ and $r_k - l_k \to L \in [c, 4c]$ for some constant $L > 0$. Then $x(r_k) = x^{l_k}(r_k - l_k) \to \bar{x}'(L)$ and $x(l_k) = x^{l_k}(0) \to \bar{x}'(0)$ by the uniform convergence of $x^{l_k}(\cdot)$ on the compact set $[0, 4c]$ and continuity of the limiting trajectory $\bar{x}'(\cdot)$. The continuity of $F(\cdot)$ then implies $F(\bar{x}'(L)) \geq F(\bar{x}'(0))$ because by assumption $F(x(r_k)) \geq F(x(l_k))$ for all k, which implies that $\bar{x}'(0)$ is stationary by Corollary 3.1. Of course, by our construction (23) of the sequences h_k and h_k', the point $\bar{x}'(0) \in A_{\rho+\delta} \setminus A_{\rho+\delta_2}$, which must not be stationary by Assumption B. This contradiction yields our desired claim (22). $\square$

With these three lemmas, we are finally ready to prove Theorem 5 proper. If x_k has a non-stationary cluster point $x_\infty \notin X^\star$, then so too does the interpolated trajectory $x(t)$ as $t \to \infty$ (recall definition (12)). Thus, suppose for the sake of contradiction that the interpolated trajectory $x(t)$ has a non-stationary cluster point $x_\infty \notin X^\star$ as $t \to \infty$. Let $\{h_k\}_{k=1}^{\infty}$, $h_k \to \infty$, be a sequence such that $x(h_k) \to x_\infty$. Using the functional convergence guaranteed by Theorem 3, some subsequence of $x^{h_k}(\cdot)$ converges to a trajectory of the differential inclusion (9) beginning from the point x_∞, which we denote by $\bar{x}(\cdot)$ (w.l.o.g. we assume for simplicity that $x^{h_k}(\cdot) \to \bar{x}(\cdot)$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$). Using the Lyapunov inequality of Theorem 4, the continuity of $\bar{x}$, and the outer-semicontinuity of $\partial f(\cdot) + \partial\varphi(\cdot) + \mathcal{N}_X(\cdot)$, we know that for all $t > 0$, $F(\bar{x}(t)) < F(x_\infty)$. Using the uniform convergence of $x^{h_k}(\cdot)$ to $\bar{x}(\cdot)$ on $[0, t]$, we see that for large enough k, $F(x(h_k + t)) < \frac{1}{2}(F(x_\infty) + F(\bar{x}(t)))$. This contradicts Lemma 3.18, as $F(x(h_k+t))$ has cluster points distinct from $F(x_\infty) = \lim_{k\to\infty} F(x(h_k))$. Finally, that $F(x_k)$ converges is immediate from Lemma 3.18.

4 Experiments

The convergence results in the preceding sections show that both the stochastic prox-linear method (5) and the stochsatic subgradient (SGD) method (6) converge to stationary points of the objective, and moreover, they have the same limiting differential inclusion. While suggestive that these methods may be effective in practice, the limiting results provide somewhat limited guidance for application of the methods. To that end, in this section we present a few experimental results explicating the performance of the methods, comparing their performance to the baseline prox-linear method (4) of Drusvyatskiy and Lewis [12, Section 5]. This deterministic provides the concrete convergence guarantee that (roughly) after $1/\epsilon^2$ iterations, the method can output an ϵ-approximate stationary point, that is, a point $\widehat{x}$ such that there exists x_0 with $\|\widehat{x} - x_0\| \le \epsilon$ and $\min\{\|g\| : g \in \partial f(x_0)\} \le \epsilon$. These comparisons provide us a somewhat better understanding of the practical advantages and disadvantages of the stochastic methods we analyze.

We consider the following problem. We have observations $b_i = \langle a_i, x^\star \rangle^2$, $i = 1, \ldots, n$, for an unknown vector $x^\star \in \mathbb{R}^d$, and we wish to find $x^\star$. This is a quadratic system of equations, which arises (for example) in phase retrieval problems in imaging science as well as in a number of combinatorial problems [8, 7]. The natural exact penalty form of this system of equations yields the minimization problem

$$\underset{x}{\text{minimize}} \;\; f(x) := \frac{1}{n} \sum_{i=1}^{n} |\langle a_i, x \rangle^2 - b_i|, \tag{24}$$

which is certainly of the form (1) with the function $h(t) = |t|$ and $c_i(x) = \langle a_i, x \rangle^2 - b_i$, so we may take the sample space $\mathcal{S} = \{1, \ldots, n\}$. In somewhat more general noise models, we may also assume we observe $b_i = \langle a_i, x^\star \rangle^2 + \xi_i$ for some noise sequence ξ_i; in this case the problem (24) is a natural robust analogue of the typical phase retrieval problem, which uses the smooth objective $(\langle a_i, x \rangle^2 - b_i)^2$, as is natural under Gaussian noise assumptions [31]. While there are a number of specialized procedures for solving such quadratic equations [8], we view problem (24) as a natural candidate for exploration of our algorithms' performance.

The stochastic prox-linear update (5) is reasonably straightforward to compute for the problem (24). Indeed, as $\nabla_x(\langle a_i, x \rangle^2 - b_i) = 2 \langle a_i, x \rangle a_i$, by rescaling by the stepsize α_k we may simplify the problem to minimizing $|b + \langle a, x \rangle| + \frac{1}{2} \|x - x_0\|^2$ for some scalar b and vectors $a, x_0 \in \mathbb{R}^d$. A standard Lagrangian calculation then shows that

$$\underset{x}{\text{argmin}} \left\{ |b + \langle a, x \rangle| + \frac{1}{2} \|x - x_0\|^2 \right\} = x_0 - \pi(\lambda) a \;\; \text{where } \lambda = \frac{\langle x_0, a \rangle + b}{\|a\|^2}$$

and $\pi(\cdot)$ is the projection of its argument into the interval $[-1, 1]$. Each iteration k of the deterministic prox-linear method [12] requires solving the quadratic program

$$x_{k+1} = \underset{x}{\text{argmin}} \left\{ \frac{1}{n} \sum_{i=1}^{n} |\langle a_i, x_k \rangle^2 + 2 \langle a_i, x_k \rangle \langle a_i, x - x_k \rangle - b_i| + \frac{1}{2\alpha} \|x - x_k\|^2 \right\}, \tag{25}$$

which we perform using Mosek via the `Convex.jl` package in `Julia` [33].

Before we present our results, we describe our choices for all parameters in our experiments. In each experiment, we let $n = 500$ and $d = 50$, and we choose $x^\star$ uniformly from the unit sphere $\mathbb{S}^{d-1}$. The noise variables ξ_i are i.i.d. Laplacian random variables with mean 0 and scale parameter σ, which we vary in our experiments. We construct the design matrix $A \in \mathbb{R}^{n \times d}$, $A = [a_1 \; \cdots \; a_n]^T$, where each row is a measurement vector a_i, as follows: we choose $U \in \mathbb{R}^{n \times d}$ uniformly from the

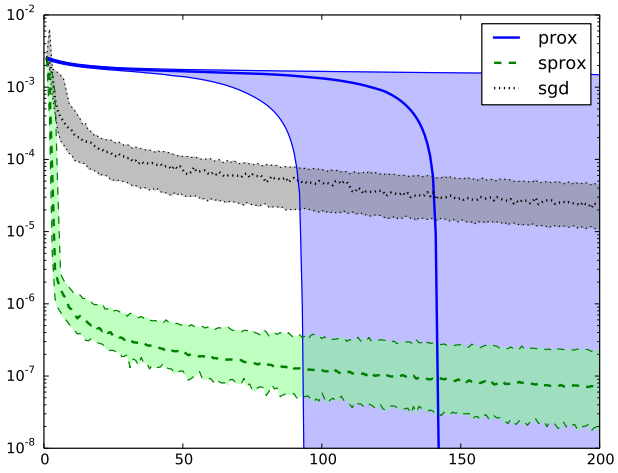

(a) Noiseless

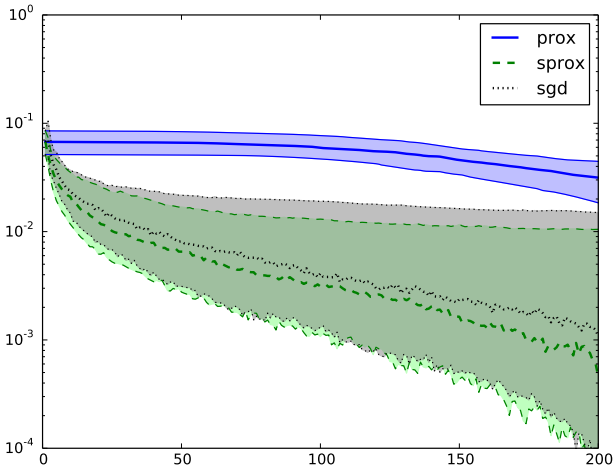

(b) Noisy

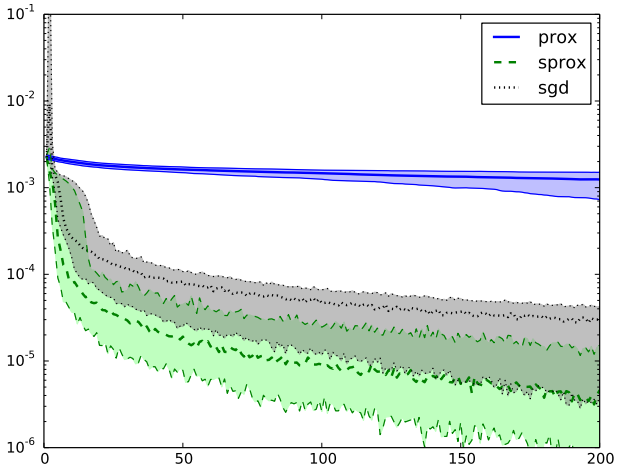

(c) Corrupted

Figure 1. Experiments with well-conditioned A matrices. In each plot, the vertical axis shows $f(x_k) - f(x^\star)$ and the horizontal axis is iteration count for the prox-linear iteration (25) or the k/n-th iteration (i.e. number of passes through the data) for the stochastic methods. The key is as follows: prox is the prox-linear iteration (25), sprox is the stochastic prox-linear method (5), and sgd is the stochastic subgradient method (6). (a) Comparison of methods with no noise, i.e. $\xi_i = 0$. (b) Comparison of methods with noise ξ_i i.i.d. Laplacian with scale $\sigma = 1$. (c) Comparison of methods with proposition $p = .1$ of observations b_i corrupted arbitrarily.

orthogonal matrices in $\mathbb{R}^{n\times d}$, i.e., $U^T U = I_{d\times d}$. We then make one of two choices for A, the first of which controls the condition number of A and the second the regularity of the norms of the rows a_i. In the former case, we set $A = UR$, where $R \in \mathrm{diag}(\mathbb{R}^d) \subset \mathbb{R}^{d\times d}$ is diagonal with linearly spaced diagonal elements in $[1, \kappa]$, so that $\kappa \geq 1$ gives the condition number of A. In the latter case, we set $A = RU$, where $R \in \mathrm{diag}(\mathbb{R}^n) \subset \mathbb{R}^{n\times n}$ is again diagonal with linearly spaced elements in $[1, \kappa]$. Finally, in each of our experiments, we set the stepsize for the stochastic methods as $\alpha_k = \alpha_0 k^{-\beta}$, where $\alpha_0 > 0$ is the initial stepsize and $\beta \in (\frac{1}{2}, 1)$ governs the rate of decrease in stepsize. We present three experiments in more detail in the coming subsections: (i) basic performance of the algorithms, (ii) the role of conditioning in the data matrix A, and (iii) an analysis of stepsize sensitivity for the different stochastic methods, that is, an exploration of the effects of the choices of α_0 and β in the stepsize choice.

4.1 Performance for well-conditioned problems

In our first group of experiments, we investigate the performance of the three algorithms under noiseless and noisy observational situations. In each of these experiments, we set the condition

number $\kappa = \kappa(A) = 1$. We consider three experimental settings to compare the procedures: in the first, we have noiseless observations $b_i = \langle a_i, x^\star \rangle^2$; in the second, we set $b_i = \langle a_i, x^\star \rangle^2 + \xi_i$ where ξ_i are Laplacian with scale $\sigma = 1$; and in the third, we again have noiseless observations $b_i = \langle a_i, x^\star \rangle^2$, but for a fraction $p = .1$ of the observations, we replace b_i with an independent $\mathsf{N}(0, 25)$ random variable, so that $n/10$ of the observations provide no information. Within each experimental setting, we perform $T = 100$ independent tests, and in each individual test we allow the stochastic methods to perform $N = 200n$ iterations (so approximately 200 loops over the data). For the deterministic prox-linear method (25), we allow 200 iterations. Each deterministic iteration is certainly more expensive than n (sub)gradient steps or stochastic prox-linear steps, but it provides a useful benchmark for comparison. The stochastic methods additionally require specification of the initial stepsize α_0 and power β for $\alpha_k = \alpha_0 k^{-\beta}$, and to choose this, we let $\alpha_0 \in \{1, 10, 10^2, 10^3\}$ and $\beta \in \{.6, .7, .8, .9\}$, perform $3n$ steps of the stochastic method with each potential pair (α_0, β), and then perform the full $N = 200n$ iterations with the best performing pair. We measure performance of the methods within each test by plotting the gap $f(x_k) - f(x^\star)$, where we approximate $x^\star$ by taking the best iterate x_k produced by any of the stochastic and deterministic methods. While the problem is non-convex and thus may have spurious local minima, these gaps provide a useful quantification of (relative) algorithm performance.

We summarize our experimental results in Figure 1. In each plot, we plot the median of the excess gap $f(x_k) - f(x^\star)$ as well as its 10% and 90% confidence intervals over our $T = 100$ tests. In order to compare the methods, the horizontal axis scales as iteration k divided by n for the stochastic methods and as iteration for the deterministic method (25). Each of the three methods is convergent in these experiments, and the stochastic methods exhibit fast convergence to reasonably accurate (say $\epsilon \approx 10^{-4}$) solutions after a few passes through the data. Eventually (though we do not always plot such results) the deterministic prox-linear algorithm achieves substantially better accuracy, though its progress is often slower. This corroborates substantial experience from the convex case with stochastic methods [c.f. 24, 14]—deterministic methods eventually find very high-quality solutions after many iterations, but stochastic methods quickly identify reasonably good solutions.

There are differences in behavior for the different methods, which we can heuristically explain. In Fig. 1(a), the stochastic prox-linear method (5) converges substantially more quickly than SGD. Intuitively, we expect this behavior because each data point (a_i, b_i) should have $\langle a_i, x \rangle^2 = b_i$ exactly, and the precise stepping of the prox-linear method achieves this more easily. In Fig. 1(b), where $b_i = \langle a_i, x^\star \rangle^2 + \xi_i$ the two methods have similar behavior; in this case, the population expectation $f_{\rm pop}(x) = \mathbb{E}[|b - \langle a, x^\star \rangle^2|]$ is smooth, because the noise ξ has a density, so gradient methods are likely to be reasonably effective. Moreover, with probability 1 we have $\langle a_i, x^\star \rangle^2 \neq b_i$, so that the precision of the prox-linear step is unnecessary. Finally, Fig. 1(c) shows that the methods are robust to corruption, but because we have $\langle a_i, x^\star \rangle^2 = b_i$ for the majority of $i \in \{1, \ldots, n\}$, there is still benefit to using the more exact (stochastic) prox-linear iteration. We note in passing that the gap in function values $f(x_k)$ between the stochastic prox-linear method and SGD is statistically significantly positive at the $p = 10^{-2}$ level for iterations $k = 1, \ldots, 20$, and that at each iteration k, the prox-linear method outperforms SGD for at least 77 of the $T = 100$ experiments (which is statistically significant for rejecting the hypothesis that each is equally likely to achieve lower objective value than the other at level $p = 10^{-6}$).

4.2 Problem conditioning and observation irregularity

In our second set of experiments, we briefly investigate conditioning of the problem (24) by modifying the condition number $\kappa = \kappa(A)$ of the measurement matrix $A \in \mathbb{R}^{n \times d}$ or by modifying the

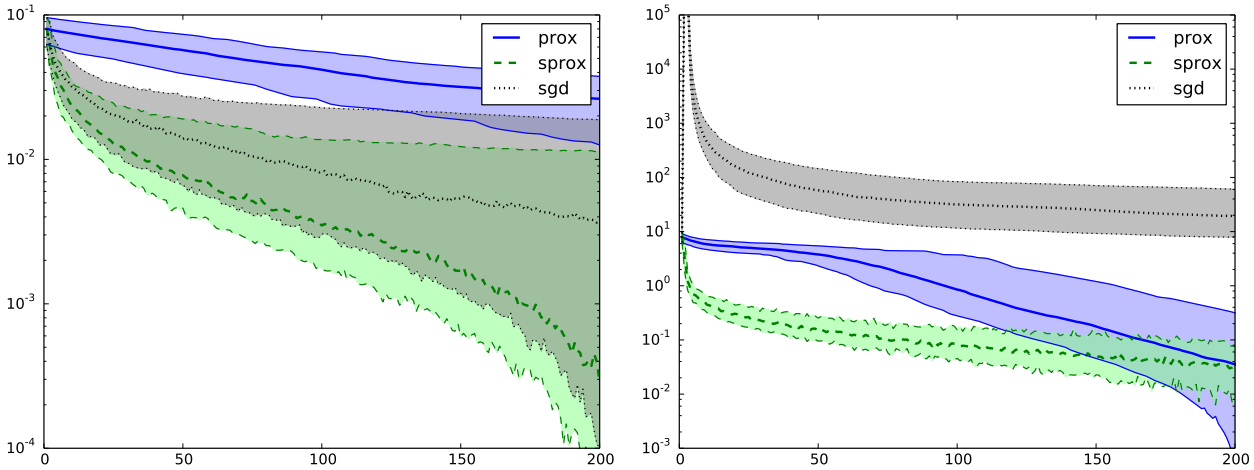

(a) $A = UR$, $\kappa(A) = 10$ (b) $A = UR$, $\kappa(A) = 100$

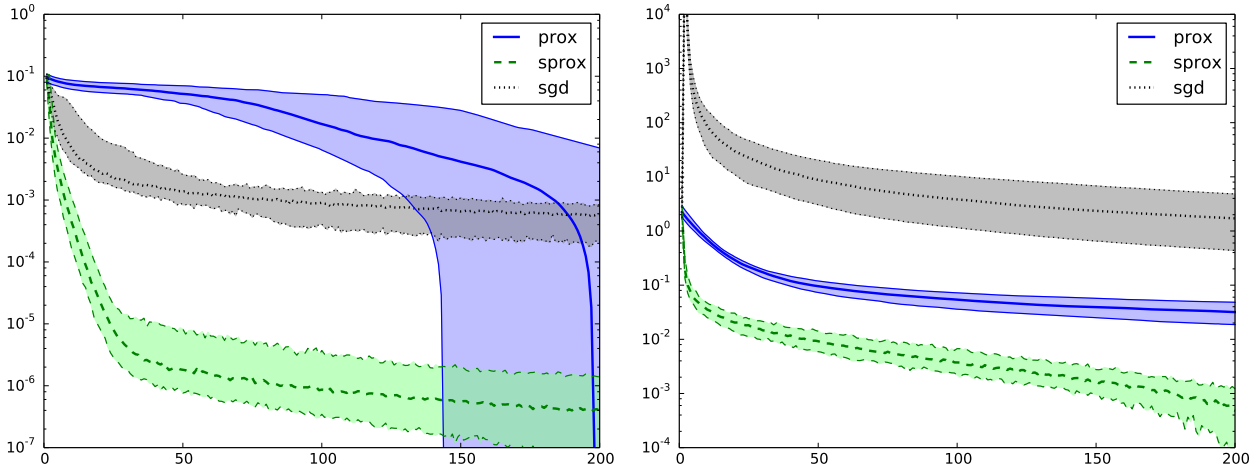

(c) $A = RU$, $b_i = \langle a_i, x^\star \rangle^2$ (d) $A = RU$, $b_i = \langle a_i, x^\star \rangle^2 + \xi_i$

Figure 2. Experiments with A matrices of varying condition number and irregularity in row norms.

relative norms of the rows $\|a_i\|$ of A. In each of the experiments, we choose the initial stepsize α_0 and power β in $\alpha_k = \alpha_0 k^{-\beta}$ using the same heuristic as the previous experiment for the stochastic methods (by considering a grid of possible values and selecting the best after $3n$ iterations). We present four experiments, whose results we summarize in Figure 2. As in the previous experiments, we plot the gaps $f(x_k) - f(x^\star)$ versus iteration k (for the deterministic prox-linear method) and versus iteration k/n for the stochastic methods. In the first two, we use observations $b_i = \langle a_i, x^\star \rangle^2 + \xi_i$, where the noise variables are i.i.d. Laplacian with scale $\sigma = 1$, and we set $A = UR$ where R is diagonal, in the first (Fig. 2(a)) scaling between 1 and $\kappa = 10$ and in the second (Fig. 2(b)) scaling between 1 and $\kappa = 100$. Each method's performance degrades as the condition number $\kappa = \kappa(A)$ increases, as one would expect. The performance of SGD degrades substantially more quickly with the conditioning of the matrix A, in spite of the fact that noisy observations improve its performance relative to the other methods (in the case $\sigma = 0$, SGD's relative performance is worse).

In the second two experiments, we set $A = RU$, where R is diagonal with entries linearly spaced in $[1, \kappa]$ for $\kappa = 10$, so that the norms $\|a_i\|$ are irregular (varying by approximately a factor

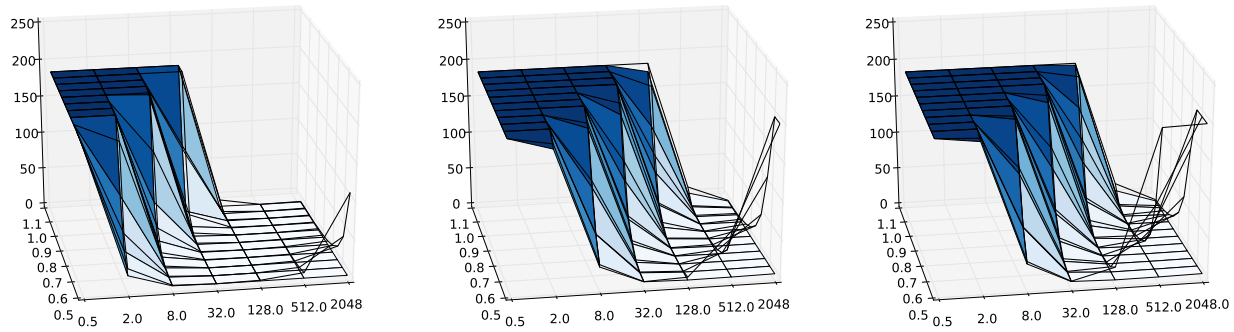

(a) $\epsilon = 5 \cdot 10^{-3}$ (b) $\epsilon = 10^{-3}$ (c) $\epsilon = 5 \cdot 10^{-4}$

Figure 3. Iterations to achieve ϵ-accuracy for the stochastic prox-linear and subgradient algorithms.

of $\kappa = 10$). In the first of the experiments (Fig. 2(c)), we set the observations $b_i = \langle a_i, x^\star \rangle^2$ with no noise, while in the second (Fig. 2(d)) we set $b_i = \langle a_i, x^\star \rangle^2 + \xi_i$ for ξ_i i.i.d. Laplacian with scale $\sigma = 1$. In both cases, the stochastic prox-linear method has better performance—this is to be expected, because its more exact updates involving the linearization $h(c(x_k; s) + \nabla c(x_k; s)^T (x - x_k); s)$ are more robust to scaling of $\|a_i\|$. As we explore more carefully in the next set of experiments, one implication of these results is that the robustness and stability of the stochastic prox-linear algorithm with respect to problem conditioning is reasonably good, while the behavior of stochastic gradient methods can be quite sensitive to conditioning behavior of the design matrix A.

4.3 Robustness of stochastic methods to stepsize

In our final experiment, we investigate the effects of stepsize parameters for the behavior of our stochastic methods. For stepsizes $\alpha_k = \alpha_0 k^{-\beta}$, the stochastic methods require specification of both the parameter α_0 and β, so it is of interest to investigate the robustness of the stochastic prox-linear method and SGD to various settings of α_0 and β. In each of these experiments, we set the condition number $\kappa(A) = 1$ and have no noise, i.e. $b_i = \langle a_i, x^\star \rangle^2$. We vary the initial stepsize $\alpha_0 \in \{2^{-1}, 2^1, 2^3, \ldots, 2^{11}\}$ and the power $\beta \in \{0.5, 0.55, 0.6, \ldots, 1\}$. In this experiment, we have $f(x^\star) = 0$, and we investigate the number of iterations

$$K(\epsilon) := \inf \{k \in \mathbb{N} \mid f(x_k) \le \epsilon\}$$

required to find and ϵ-optimal solution. (In our experiments, the stochastic methods always find such a solution eventually.) We perform $T = 250$ tests for each setting of the pairs α_0, β, and in each test, we run both the stochastic prox-linear method (5) or SGD (6) for $k = 200n$ iterations, setting $K(\epsilon) = 200n$ if no function value $f(x_k)$ satisfies $f(x_k) \le \epsilon$ within this window.

Figure 3 illustrates the results of these experiments, where the vertical axis gives the median time to ϵ-accuracy over all the $T = 250$ tests versus the parameter α_0 and β, which correspond to the two horizontal axes. Each plot corresponds to a different value of ϵ (which both methods achieve, as demonstrated by Fig. 1(a)). The solid white-to-blue surface, with thin lines, corresponds to the iteration counts for the stochastic prox-linear method (5); the transparent surface with thicker lines corresponds to the iteration counts for the stochastic subgradient method (6). Figure 3 shows that the stochastic prox-linear algorithm consistently has comparable or better performance than SGD for the same choices of parameters α_0, β. The more salient aspect of the figures is, however,

that the stochastic prox-linear method is more robust to stepsize (mis-)specification than is SGD. Indeed, as made most apparent in Fig. 3(c), the range of stepsizes yielding good performance for SGD is a relatively narrow valley, while the prox-linear method enjoys reasonable performance for essentially all large enough stepsizes α_0 and slow enough rates of decrease in the stepsize, that is, when β is near $\frac{1}{2}$ and $\alpha_0 > 30$ or so. This behavior is perhaps expected: the iterations (5) guard more carefully against wild swings that result from aggressive stepsize choices, yielding more robust convergence. Thus, in spite of the identical limiting differential inclusion that both methods attain, as described by Theorems 1 and 3, it is likely easier to tune the stepsizes to achieve good behavior for the stochastic prox-linear algorithm than for the stochastic sub-gradient method.

A Proof of Theorem 2

To prove the theorem, we require a few pieces of additional notation. Define the indices

$$k_<(t) = \max\{k \in \mathbb{N} : t_k \le t\} \;\text{ and }\; k_>(t) = \min\{k \in \mathbb{N} : t_k \ge t\} \;\text{ for } t \in \mathbb{R}_+.$$

For $t \in \mathbb{R}_+$, let $x_t(\cdot)$ denote the solution to the differential equation

$$\dot{x}_t(\cdot) = y(\cdot) \;\text{ and }\; x_t(t) = x(t), \tag{26}$$

which is evidently $x_t(\tau) = x(t) + \int_t^\tau y(u)du$. We divide the proof into two main parts: first, we show that the function family $x^\tau(\cdot)$ is relatively compact in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ using the Arzelà-Ascoli theorem. In the second part, we apply a few functional-analytic results on weak convergence to establish that $x^\tau(\cdot)$ has the limiting differential inclusion properties claimed in the theorem.

Part I: Relative compactness We begin with a lemma that shows an equivalence between the time-shifted and interpolated (12) sequence $x^t(\cdot) = x(t + \cdot)$ and the differential sequence $x_t(\cdot)$ as $t \to \infty$, showing that the noise and discretization effects in the iteration (11) are essentially negligible. As we show presently, the relative compactness of $x_t(\cdot)$ is reasonably straightforward, so this equivalence simplifies our arguments.

Lemma A.1. *For any $T \ge 0$, we have*

$$\lim_{t\to\infty} \sup_{\tau\in[t,t+T]} \|x(\tau) - x_t(\tau)\| = 0.$$

Proof Fix $\tau \in [t, t+T]$ and let $p = k_>(t)$ and $q = k_<(\tau)$. Then by definition of x_t we have

$$\begin{aligned}
x_t(\tau) &= x(t) + \int_t^{t_p} y(u)du + \int_{t_p}^{t_q} y(u)du + \int_{t_q}^{\tau} y(u)du \\
&= x(t_p) + \sum_{i=p}^{q-1} \alpha_i y_i + \left(x(t) + \int_t^{t_p} y(u)du - x(t_p)\right) + \int_{t_q}^{\tau} y(u)du \\
&= x(t_q) - \sum_{i=p}^{q-1} \alpha_i \xi_i + \left(x(t) + \int_t^{t_p} y(u)du - x(t_p)\right) + \int_{t_q}^{\tau} y(u)du \\
&= x(\tau) - \sum_{i=p}^{q-1} \alpha_i \xi_i + \left(x(t) + \int_t^{t_p} y(u)du - x(t_p)\right) + \left(x(t_q) + \int_{t_q}^{\tau} y(u)du - x(\tau)\right).
\end{aligned} \tag{27}$$

Taking norms and appling the triangle inequality, we have

$$\left\|x(t) + \int_t^{t_p} y(u)du - x(t_p)\right\| \le \|x(t_p) - x(t_{p-1})\| + \int_t^{t_p} \|y(u)\| \, du \le \alpha_{p-1} \left(\|\xi_{p-1}\| + 2\|y_{p-1}\|\right),$$

and similarly

$$\left\|x(t_q) + \int_{t_q}^{\tau} y(u)du - x(\tau)\right\| \le \|x(t_{q+1}) - x(t_q)\| + \int_{t_q}^{\tau} \|y(u)\| \, du \le \alpha_q \left(\|\xi_q\| + 2\|y_q\|\right).$$

The convergence of the sum $\sum_{k=1}^{\infty} \alpha_k \xi_k$ implies the uniform guarantees

$$\lim_{p\to\infty} \sup_{q \ge p} \left\|\sum_{i=p}^{q-1} \alpha_i \xi_i\right\| = 0 \quad \text{and} \quad \lim_{p\to\infty} \sup_{q\ge p} \alpha_q(\|\xi_q\| + \|y_q\| = 0.$$

Substituting these into the bound (27) gives the result. □

Now we may show the relative compactness of the function sequence $x^{t_k}(\cdot)$ for any sequence t_k. Using that $\sup_t \|y(t)\| = \sup_k \|y_k\| < \infty$ by assumption, we have that $\{x_t(\cdot)\}_{t\in\mathbb{R}_+}$ is a family of uniformly equicontinuous (even Lipschitz) and pointwise bounded functions. The Arzelà-Ascoli theorem implies it is therefore relatively compact in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$. In turn, this implies that the family of shifted interpolant functions $\{x^t(\cdot)\}_{t\in\mathbb{R}_+} \in \mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ is also relatively compact. Indeed, pick an arbitrary sequence $\{t_k\}_{k=1}^{\infty}$; then one of the following two cases must occur.

(i) The sequence $\{t_k\}_{k=1}^{\infty}$ has a bounded subsequence. Pick any cluster point t of t_k (w.l.o.g. we may assume $t_k \to t$). Then by the uniform (even Lipschitz) continuity of $x(\cdot)$ on $[0, t+T]$ for any $T \ge 0$, we have that $x^{t_k}(\cdot) \to x^t(\cdot)$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$.

(ii) The sequence $\{t_k\}_{k=1}^{\infty}$ satisfies $t_k \to \infty$. Since $x_{t_k}(\cdot)$ is relatively compact in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$, it has a convergence subsequence $x_{t_{k_n}}(\cdot) \to \overline{x}(\cdot)$ for some limit function $\overline{x}(\cdot)$. Lemma A.1 shows that $x^{t_{k_n}}(\cdot) \to \overline{x}(\cdot)$ as well.

These cases combined yield that $\{x^t(\cdot)\}_{t\in\mathbb{R}_+}$ is relatively compact.

Part II: Limiting differential inclusion We now establish the remainder of the theorem. Following our notational conventions, we define the shifts $y^t(\cdot) = y(t + \cdot)$. Let $\{u_k\}_{k\in\mathbb{N}}$ satisfy $u_k \to \infty$ and fix $T > 0$, and w.l.o.g. assume that $x^{u_k}(\cdot) \to \overline{x}(\cdot)$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$ for some continuous $\overline{x}(\cdot)$. Viewing $\mathcal{Y}_T := \{y^{u_k}(\tau), \tau \in [0,T]\}_{k\in\mathbb{N}}$ as a subset of the Hilbert space $L^2([0,T])$, $\mathcal{Y}_T$ is bounded and thus weakly sequentially compact by the Banach-Alaoglu theorem. Thus the sequence $\{u_k\}$ has a subsequence u_{k_n} such that $y^{u_{k_n}}(\cdot) \to \overline{y}(\cdot)$ weakly in $L^2([0,T])$ for some $\overline{y}(\cdot) \in L^2([0,T])$. Lemma A.1 also implies that the integrated sequence $x_{u_k}(u_k + \cdot) \to \overline{x}(\cdot)$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$. Using the definition (26) of $x_t(t+\cdot)$, we have for all $t, \tau \in \mathbb{R}_+$ that

$$x_t(t+\tau) = x_t(t) + \int_0^{\tau} y^t(u)du.$$

Substituting $t = u_{k_n}$, and taking limits as $k_n \to \infty$ on both sides of this equality, we obtain that for $\tau \le T$ we have

$$\overline{x}(\tau) = \overline{x}(0) + \int_0^{\tau} \overline{y}(u)du \tag{28}$$

by the weak convergence of $y^{u_{k_n}}$. It remains to show that $\overline{y}(u) \in H(\overline{x}(u))$ for (almost) all $u \leq T$.

As the set $\mathcal{Y}_T \subset L^2([0,T])$ is bounded, the Banach-Saks theorem implies the existence of a further subsequence $\{k_{n_j}\}_{j\in\mathbb{N}}$, which to simplify notation we now assume without loss of generality is simply the original sequence k, such that

$$\frac{1}{N}\sum_{k=1}^{N} y^{u_k}(\cdot) \to \overline{y}(\cdot)$$

in $L^2([0,T])$ as well as almost everywhere. (We could replace the original sequence u_k with $u_{k_{n_j}}$ and repeat, *mutatis mutandis*, our preceding argument.) For $t \in \mathbb{R}_+$, define the left values $l_k(t) := k_<(u_k + t)$. Then we have by definition of the interpolated sequence $y(\cdot)$ and the sequence of mappings $g_k : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ that

$$y^{u_k}(\tau) \in g_{l_k(\tau)}(x(t_{l_k(\tau)}))$$

where we recall that $t_l = \sum_{k=1}^{l} \alpha_k$. By the definition $x^t(\tau) = x(t+\tau)$, the triangle inequality, and the left indices $l_k(\tau)$, for any fixed τ we find that

$$\begin{aligned}
\left\|x(t_{l_k(\tau)}) - \overline{x}(\tau)\right\| &\leq \left\|x(t_{l_k(\tau)}) - x(u_k+\tau)\right\| + \left\|x^{u_k}(\tau) - \overline{x}(\tau)\right\| \\
&\leq \left\|x(t_{l_k(\tau)+1}) - x(t_{l_k(\tau)})\right\| + \left\|x^{u_k}(\tau) - \overline{x}(\tau)\right\| \\
&\leq \alpha_{l_k(\tau)}\left\|\xi_{l_k(\tau)}\right\| + \alpha_{l_k(\tau)}\left\|y_{l_k(\tau)}\right\| + \left\|x^{u_k}(\tau) - \overline{x}(\tau)\right\|,
\end{aligned}$$

and this quantity converges to zero as $k \to \infty$ because $y(\cdot)$ is bounded by condition (i) of the theorem, $\alpha_k \left\|\xi_k\right\| \to 0$ by the convergence condition (iii), and $x^{u_k}(\tau) \to \overline{x}(\tau)$ by the assumption that $x^{u_k} \to \overline{x}$ in $\mathcal{C}(\mathbb{R}_+, \mathbb{R}^d)$. That is, $x(t_{l_k(\tau)}) - \overline{x}(\tau) \to 0$ for all τ as $k \to \infty$.

Lastly, we apply item (iv) in the conditions for the theorem. Using $y^{u_k}(\tau) \in g_{l_k(\tau)}(x(t_{l_k(\tau)}))$, we have

$$\begin{aligned}
\operatorname{dist}(\overline{y}(\tau), H(\overline{x}(\tau))) &\leq \underbrace{\left\|\frac{1}{N}\sum_{k=1}^{N} y^{u_k}(\tau) - \overline{y}(\tau)\right\|}_{\to 0} + \operatorname{dist}\left(\frac{1}{N}\sum_{k=1}^{N} y^{u_k}(\tau), H(\overline{x}(\tau))\right) \\
&\leq \operatorname{dist}\left(\frac{1}{N}\sum_{k=1}^{N} g_{l_k(\tau)}(x(t_{l_k(\tau)})), H(\overline{x}(\tau))\right) + o(1) \to 0
\end{aligned}$$

as $N \to \infty$ by item (iv). As $H(\overline{x}(\tau))$ is closed, we obtain $\overline{y}(\tau) \in H(\overline{x}(\tau))$ for almost every $\tau \in [0,T]$, as desired. Modifying $\overline{y}(\tau)$ on a suitable null set and recognizing that T is arbitrary then yields the theorem.

B Technical Proofs and Results

B.1 An integral distance inequality

Lemma B.1. *Let $a, b : \mathcal{S} \to \mathbb{R}^d$ be measurable and $C, D : \mathcal{S} \rightrightarrows \mathbb{R}^d$ be measurable and closed-convex-valued. Then for any measure μ on $\mathcal{S}$,*

$$\operatorname{dist}\left(\int (a(s)+b(s))d\mu(s), \int (C(s)+D(s))d\mu(s)\right) \leq \int \left[\operatorname{dist}(a(s), C(s)) + \operatorname{dist}(b(s), D(s))\right] d\mu(s).$$

Proof First, we note that for any vector $v \in \mathbb{R}^d$, the mapping to the projection $s \mapsto \pi_{C(s)}(v)$ is measurable [28, Lemma 1] and $v \mapsto \pi_{C(s)}(v)$ is Lipschitz continuous as it is a projection. Thus, by a standard approximation argument with simple functions, for any measurable $c : \mathcal{S} \to \mathbb{R}^d$, the function $s \mapsto \pi_{C(s)}(c(s))$ is measurable. As a consequence, we have the following string of inequalities:

$$\begin{aligned}
&\operatorname{dist}\left(\int (a+b)d\mu, \int (C+D)d\mu\right) \\
&= \inf\left\{\left\|\int (a+b)d\mu - \int (c+d)d\mu\right\| : c(s) \in C(s), d(s) \in D(s), c, d \text{ measurable}\right\} \\
&\le \inf\left\{\int \|(a+b)-(c+d)\| \, d\mu : c(s) \in C(s), d(s) \in D(s), c, d \text{ measurable}\right\} \\
&\le \inf\left\{\int \|a-c\| \, d\mu + \int \|b-d\| \, d\mu : c(s) \in C(s), d(s) \in D(s), c, d \text{ measurable}\right\} \\
&= \int \left\|a(s) - \pi_{C(s)}(a(s))\right\| d\mu(s) + \int \left\|b(s) - \pi_{D(s)}(b(s))\right\| d\mu(s),
\end{aligned}$$

where the final equality follows because $\pi_{C(s)}(v) \in C(s)$ and $s \mapsto \pi_{C(s)}(a(s))$ is measurable. The integrands in the above are simply $\operatorname{dist}(a(s), C(s))$ and $\operatorname{dist}(b(s), D(s))$, which gives the result. $\square$

B.2 Proof of Lemma 3.6

Let $B \ge \|x\| + 1$ and $\gamma(s) \equiv \gamma(B; s)$ and $\beta(s) \equiv \beta(B; s)$ be as in Assumption A. It is an essentially immediate and standard calculation [12, 13] that if we let $f(x, s) = h(c(x; s); s)$ for shorthand, then $\partial f(x, s) = \nabla c(x; s) \partial h(c(x; s); s)$. Thus, letting $f'(x, s; v)$ be the (Clarke) directional derivative of $f(\cdot, s)$ in the direction v and $h'(c; s; v)$ be the directional derivative of $h(\cdot; s)$ in the direction v at the point c, we have $f'(x, s; v) = \sup_{g \in \partial f(x,s)} \langle g, v \rangle = h'(c; s; \nabla c(x; s)^T v)$.

We now integrate this expression. For $u \in \mathbb{R}^d$ and $t < 1/\|u\|$, we have

$$\begin{aligned}
f(x+tu) &= \int h(c(x+tu; s); s) dP(s) \\
&\in \int h(c(x; s) + t \nabla c(x; s)^T u; s) dP(s) \pm \int \frac{t^2 \gamma(s) \beta(s)}{2} \|u\|^2 \, dP(s).
\end{aligned}$$

Because $u \mapsto h(c(x; s) + t \nabla c(x; s)^T u; s)$ is a normal convex integrand [29, Ch. 14], this containment, the $\gamma(s)$-Lipschitz continuity of $h(\cdot; s)$, that there exists $x_0 \in X$ such that $\mathbb{E}[\|\nabla c(x_0; S)\|^2] < \infty$, and the fact that $\mathbb{E}[\gamma(S)\beta(S)] < \infty$ imply by the dominated convergence theorem that

$$\begin{aligned}
\frac{f(x+tu) - f(x)}{t} &= \int \frac{h(c(x; s) + t \nabla c(x; s)^T u; s) - h(c(x; s); s)}{t} dP(s) \pm \frac{t \|u\|^2}{2} \mathbb{E}[\gamma(S)\beta(S)] \\
&\to \int h'(c(x; s); s; \nabla c(x; s)^T v) dP(s) = \int f'(x, s; v) dP(s) \quad \text{as } t \to 0, u \to v.
\end{aligned}$$

That is, $f'(x; v) = \int f'(x, s; v) dP(s)$. An argument completely parallel to that of Bertsekas [3, Propositions 2.1–2.2] yields that $\partial f(x) = \int \nabla c(x; s) \partial h(c(x; s); s) dP(s)$ and that $\partial f(x)$ is compact.

Lastly, we show that $\partial f(\cdot)$ is outer semi-continuous. Because the support function of the subdifferential $\partial f(x)$ is the directional derivative of f, the outer semi-continuity of ∂f is equivalent to

$$\limsup_{k \to \infty} f'(x_k; v) \le f'(x; v) \quad \text{for all } \|v\| = 1 \text{ and } x_k \to x \in X$$

(see, e.g. [19, Propositions V.3.3.9]). Because the sets $\partial h(c(x;s);s)$ are locally bounded in a neighborhood of x (because c is smooth and h is Lipschitz), we use Corollary VI.6.2.5 of [19] to obtain

$$\lim_{k\to\infty}\left[h'(c(x_k;s);s;\nabla c(x_k)^T v) - h'(c(x_k;s);s;\nabla c(x)^T v)\right] = 0$$

and

$$\limsup_k h'(c(x_k;s);s;\nabla c(x)^T v) \le h'(c(x;s);s;\nabla c(x)^T v).$$

The Lipschitz condition on $h(\cdot;s)$ then gives a sufficient integrable envelope to apply Fatou's lemma, which yields

$$\begin{aligned}\limsup_k \int h'(c(x_k;s);s;\nabla c(x_k;s)^T v)dP(s) &\le \int \limsup_k h'(c(x_k;s);s;\nabla c(x_k;s)^T v)dP(s)\\ &\le \int h'(c(x;s);s;\nabla c(x;s)^T v)dP(s) = \mathbb{E}_P[f'(x,S;v)].\end{aligned}$$

That is, $\limsup_k f'(x_k;v) \le \mathbb{E}_P[f'(x,S;v)] = f'(x;v)$.

B.3 Proof of Lemma 3.8

That $L_\epsilon(x;s) < \infty$ for all x is immediate, as we have $G(x;s) = \partial\varphi(x) + \nabla c(x;s)\partial h(c(x;s);s)$, and subdifferentials of convex functions are compact convex sets.

We first show the upper semicontinuity of the function $L_\epsilon(\cdot;s)$. Suppose for the sake of contradiction that for some $x \in X$ and for some sequence $\{x_k\}_{k=1}^\infty \subset X$ converging to x, we have $\lim_{k\to\infty} L_\epsilon(x_k;s)$ exists and that there is some $\delta > 0$ such that

$$\lim_{k\to\infty} L_\epsilon(x_k;s) \ge L_\epsilon(x;s) + \delta.$$

Now, by definition of L_ϵ we may choose $x_k' \in X$ such that $\|x_k - x_k'\| \le \epsilon$ and subgradient vectors $p_k(s) \in \partial h(c(x_k';s);s)$ and $q_k(s) \in \partial\varphi(x_k')$ satisfying

$$L_\epsilon(x_k;s) \le \left\|\nabla c(x_k';s)p_k(s) + q_k(s)\right\| + \delta/2 \;\text{ for all } k.$$

Since the sequence $\{x_k'\} \subset X$ is bounded, it has accumulation points and we may assume w.l.o.g. that $x_k' \to x' \in X$, where x' satisfies $\|x' - x\| \le \epsilon$. By Assumption A, we certainly have $\nabla c(x_k';s) \to \nabla c(x';s)$ and $c(x_k';s) \to c(x's)$, and the outer semicontinuity of the differential for convex functions (Lemma 3.1) shows that there must be a subsequence $\{n_k\}$ satisfying $p_{n_k}(s) \to p(s) \in \partial h(c(x';s);s)$ and $q_{n_k}(s) \to q(s) \in \partial\varphi(x')$. In particular,

$$\begin{aligned}\lim_{k\to\infty} L_\epsilon(x_k;s) = \limsup_{k\to\infty} L_\epsilon(x_{n_k};s) &\le \limsup_{k\to\infty}\left\|\nabla c(x_{n_k}';s)p_{n_k}(s) + q_{n_k}(s)\right\| + \frac{\delta}{2}\\ &= \left\|\nabla c(x';s)p(s) + q(s)\right\| + \frac{\delta}{2} \le L_\epsilon(x;s) + \frac{\delta}{2},\end{aligned}$$

which is a contradiction as desired. Thus $L_\epsilon(x;s)$ is upper semicontinuous in x (and so is $L_\epsilon^2(x;s)$).

To see that $L_\epsilon(\cdot)$ is upper semicontinuous, we construct an integrable envelope for the function and then apply Fatou's lemma. Indeed, for any $B < \infty$ and x with $\|x\| \le B - 1$, using the shorthand $\beta(s) = \beta(B,s)$ and $\gamma(s) = \gamma(B,s)$ for the Lipschitz constants of ∇c and h, respectively (by Assumption A), we have

$$\|G(x;s)\| \le \left|\!\left|\!\left|\nabla c(x;s)\right|\!\right|\!\right|_{\text{op}} \|\partial h(c(x;s);s)\| + \|\partial\varphi(x)\|,$$

and by the Fenchel-Young inequality and Assumption A, we have

$$\begin{aligned}\left\|\nabla c(x;s)\right\|_{\rm op}\left\|\partial h(c(x;s);s)\right\| &\le \frac{1}{2}\left\|\nabla c(x;s)\right\|_{\rm op}^2 + \frac{1}{2}\left\|\partial h(c(x;s);s)\right\|^2 \\ &\le \frac{1}{2}\left(\left\|\nabla c(x_0;s)\right\|_{\rm op} + \beta(s)\left\|x-x_0\right\|\right)^2 + \frac{1}{2}\gamma(s)^2.\end{aligned}$$

This quantity is integrable, and we may thus apply Fatou's Lemma to obtain

$$\limsup_{y\to x} L_\epsilon(y) = \sqrt{\limsup_{y\to x}\mathbb{E}[L_\epsilon(y;S)^2]} \le \sqrt{\mathbb{E}\left[\limsup_{y\to x} L_\epsilon(y;S)^2\right]} \le \mathbb{E}\left[L_\epsilon(x;S)^2\right]^{\frac{1}{2}}.$$

B.4 Proof of Lemma 3.14

In the case that $\varphi(x) = \frac{\lambda}{2}\left\|x\right\|^2$, the stochastic update (6) becomes $x_{k+1} = \frac{1}{1+\alpha_k\lambda}x_k - \frac{\alpha_k}{1+\alpha_k\lambda}g_k$, and we have the recursion

$$\left\|x_{k+1}\right\| \le \frac{\left\|x_k\right\|}{1+\alpha_k\lambda} + \frac{\alpha_k}{1+\alpha_k\lambda}L(S_k) \le \prod_{i=1}^k (1+\alpha_i\lambda)^{-1}\left\|x_1\right\| + \sum_{i=1}^k \alpha_i L(S_i)\prod_{j=i}^k (1+\alpha_j\lambda)^{-1}.$$

Let $L_i = L(S_i)$ for shorthand and define $\xi_i = L_i - \mathbb{E}[L(S_i)]$, noting ξ_i are i.i.d. and mean zero. Assume that $\lambda = 1$ and $\mathbb{E}[L(S)^2] = 1$ without loss of generality. Defining $Z_k = \sum_{i=1}^k \alpha_i L_i \prod_{j=i}^k (1+\alpha_j\lambda)^{-1}$, so that $Z_k - \mathbb{E}[Z_k] = \sum_{i=1}^k \alpha_i\xi_i\prod_{j=i}^k(1+\alpha_j\lambda)^{-1}$ and noting that

$$\mathbb{E}[Z_{k+1}] = \frac{\mathbb{E}[Z_k] + \alpha_{k+1}}{1+\alpha_{k+1}} = \begin{cases} \le \mathbb{E}[Z_k] & \text{if } \mathbb{E}[Z_k] > 1 \\ \le 1 & \text{if } \mathbb{E}[Z_k] \le 1 \end{cases}$$

we have that $\sup_k \mathbb{E}[Z_k] < \infty$. Moreover, if we let $M_k = \sum_{i=1}^k \alpha_i \prod_{j=1}^{i-1}(1+\alpha_j)\xi_i$, then $Z_k - \mathbb{E}[Z_k] = \prod_{j=1}^k (1+\alpha_j)^{-1} M_k$, and M_k is a martingale adapted to the filtration $\mathcal{F}_k = \sigma(S_1,\ldots,S_k)$. Noting that $M_{k+1} - M_k = \alpha_{k+1}\prod_{j=1}^k(1+\alpha_j)\xi_{k+1}$, we have

$$\sum_{k=1}^\infty \frac{1}{\prod_{j=1}^k (1+\alpha_j)^2}\mathbb{E}[(M_{k+1}-M_k)^2 \mid \mathcal{F}_k] = \sum_{k=1}^\infty \alpha_{k+1}^2\mathbb{E}[\xi_{k+1}^2] < \infty.$$

Applying standard L_2-martingale convergence results (e.g. [10, Exercise 5.3.35]) gives that $M_k/\prod_{j=1}^k(1+\alpha_j) \stackrel{a.s.}{\to} 0$, and thus $Z_k \stackrel{a.s.}{\to} \mathbb{E}[Z_k]$, while certainly $\limsup_k \left\|x_k\right\| \le \limsup_k Z_k$.

B.5 Proof of Lemma 3.15

The standard first-order optimality conditions for minimization of (strongly) convex problems (e.g. the iteration (6)) immediately yield

$$\frac{1}{\alpha_k}\left\|x_k - x_{k+1}\right\|^2 + \left\langle g_k, x_{k+1}\right\rangle + \varphi(x_{k+1}) \le \left\langle g_k, x_k\right\rangle + \varphi(x_k), \tag{29}$$

and as the iterations are unconstrained, we have $x_{k+1} - x_k = -\alpha_k g_k - \alpha_k v_{k+1}$ for some $v_{k+1} \in \partial\varphi(x_{k+1})$. We first assume that the regularizer φ increases in iteration k, so $\varphi(x_k) \le \varphi(x_{k+1})$. Then inequality (29) implies that

$$\left\|x_{k+1} - x_k\right\|\left\|g_k + v_{k+1}\right\| = \frac{1}{\alpha_k}\left\|x_{k+1}-x_k\right\|^2 \le \left\langle g_k, x_k - x_{k+1}\right\rangle \le \left\|g_k\right\|\left\|x_k - x_{k+1}\right\|,$$

or $\|g_k + v_{k+1}\| \le \|g_k\|$ and we have $\|v_{k+1}\| \le 2\|g_k\|$. The local Lipschitz assumption on $h(c(x;s);s)$ guarantees $\|g_k\| \le L(1+\|x_k\|^\nu)$. Choose the constant $1 \le B < \infty$ such that if $\|y\| \ge B$ then $\|v\| \ge 4L\|y\|^{\beta-1}$ for all $v \in \partial\varphi(y)$, which must exist by our coercivity assumption, and $B^{\frac{\beta-1}{\nu}} \ge (1/\lambda)B$, where $\lambda \in (0,1]$ is the regularity parameter of φ for which $\varphi(\lambda x) \le \varphi(x)$ for all x with $\|x\| \ge B$ (recall $\beta - 1 > \nu$). Then we have that

$$4L\|x_{k+1}\|^{\beta-1} \le 2\|g_k\| \le 2L(1+\|x_k\|^\nu) \le 4L \cdot \begin{cases} \|x_k\|^\nu & \text{if } \|x_k\| \ge 1 \\ 1 & \text{otherwise} \end{cases}$$

if $\|x_{k+1}\| \ge B$. Otherwise, we have $\|x_{k+1}\| < B$, so our assumptions on B guarantee

$$\|x_{k+1}\| \le \max\left\{B, 1, \|x_k\|^{\frac{\nu}{\beta-1}}\right\} \le \max\left\{B, \|x_k\|^{\frac{\nu}{\beta-1}}\right\} \le \max\left\{B, \lambda\|x_k\|\right\}.$$

We assumed that $\varphi(x_{k+1}) \ge \varphi(x_k)$, so we see that (because of the λ term above)

$$\varphi(x_{k+1}) \le \varphi(x_k) \quad \text{or} \quad \|x_{k+1}\| \le B. \tag{30}$$

As φ is continuous (it is convex and defined on $\mathbb{R}^d$), there exists B' such that $\varphi(x) \le B'$ for all $\|x\| \le B$, and thus inequality (30) implies

$$\varphi(x_{k+1}) \le \varphi(x_k) \vee B'.$$

We may assume w.l.o.g. that $\varphi(x_1) \le B'$, because otherwise we could simply increase $B' = \varphi(x_1)$. An inductive argument implies that $\varphi(x_{k+1}) \le B'$ for all k, which implies the result.

References

[1] J.-P. Aubin and A. Cellina. *Differential Inclusions: Set-Valued Maps and Viability Theory*. Springer-Verlag, 1984.

[2] M. Benaïm, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44(1):328–348, 2005.

[3] D. P. Bertsekas. Stochastic optimization problems with nondifferentiable cost functionals. *Journal of Optimization Theory and Applications*, 12(2):218–231, 1973.

[4] V. Borkar. *Stochastic Approximation*. Cambridge University Press, 2008.

[5] J. Burke. Descent methods for composite nondifferentiable optimization problems. *Mathematical Programming*, 33:260–279, 1985.

[6] E. Candès, T. Strohmer, and V. Voroninski. PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8), 2013.

[7] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.

[8] Y. Chen and E. Candès. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Communications on Pure and Applied Mathematics*, To appear, 2015.

[9] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, September 1995.

[10] A. Dembo. Lecture notes on probability theory: Stanford statistics 310. Accessed October 1, 2016, 2016. URL `http://statweb.stanford.edu/~adembo/stat-310b/lnotes.pdf`.

[11] D. Drusvyatskiy and C. Kempton. An accelerated algorithm for minimizing convex compositions. *arXiv:1605.00125 [math.OC]*, 2016.

[12] D. Drusvyatskiy and A. Lewis. Error bounds, quadratic growth, and linear convergence of proximal methods. *arXiv:1602.06661 [math.OC]*, 2016.

[13] D. Drusvyatskiy, A. Ioffe, and A. Lewis. Nonsmooth optimization using taylor-like models: error bounds, convergence, and termination criteria. *arXiv:1610.03446 [math.OC]*, 2016. URL `https://arxiv.org/abs/1610.03446`.

[14] J. C. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.

[15] Y. M. Ermoliev. On the stochastic quasi-gradient method and stochastic quasi-Feyer sequences. *Kibernetika*, 2:72–83, 1969.

[16] R. Fletcher. A model algorithm for composite nondifferentiable optimization problems. *Mathematical Programming Study*, 17:67–76, 1982.

[17] R. Fletcher and G. A. Watson. First and second order conditions for a class of nondifferentiable optimization problems. *Mathematical Programming*, 18:291–307, 1980.

[18] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, second edition, 2009.

[19] J. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms I & II*. Springer, New York, 1993.

[20] M. Kunze. *Non-Smooth Dynamical Systems*. Lecture Notes in Mathematics. Springer, 2000.

[21] H. J. Kushner and G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, second edition, 2003.

[22] A. S. Lewis and S. J. Wright. A proximal method for composite minimization. *arXiv:0812.0423 [math.OC]*, 2008. URL `http://arxiv.org/abs/0812.0423`.

[23] L. Ljung. Analysis of recursive stochastic algorithms. *IEEE Transactions on Automatic Control*, 22(4):551–575, 1977.

[24] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

[25] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Kluwer Academic Publishers, 2004.

[26] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.

[27] H. Robbins and D. Siegmund. A convergence theorem for non-negative almost supermartingales and some applications. In *Optimizing Methods in Statistics*, pages 233–257. Academic Press, New York, 1971.

[28] R. T. Rockafellar. Measurable dependence of convex sets and functions on parameters. *Journal of Mathematical Analysis and Applications*, 28:4–25, 1969.

[29] R. T. Rockafellar and R. J. B. Wets. *Variational Analysis*. Springer, New York, 1998.

[30] Y. Schechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev. Phase retrieval with application to optical imaging. *IEEE Signal Processing Magazine*, pages 87–109, May 2015.

[31] J. Sun, Q. Qu, and J. Wright. When are nonconvex problems not scary? *arXiv:1510.06096 [math.OC]*, 2015.

[32] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc B.*, 58 (1):267–288, 1996.

[33] M. Udell, K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd. Convex optimization in Julia. In *First Workshop for High Performance Technical Computing in Dynamic Languages*, pages 18–28. IEEE, 2014.